

Calibrated Incentive Contracts

Sylvain Chassang^{*†}

Princeton University

December 12, 2011

Abstract

This paper studies a dynamic agency problem which includes limited liability, moral hazard and adverse selection. The paper develops a robust approach to dynamic contracting based on calibrating the payoffs that would have been delivered by simple benchmark contracts that are attractive but infeasible, due to limited liability constraints. The resulting dynamic contracts are detail-free and satisfy robust performance bounds independently of the underlying process for returns, which need not be i.i.d. or even ergodic.

1 Introduction

This paper considers a dynamic agency problem in which a principal hires an agent to make investment decisions on her behalf.¹ The contracting environment includes limited liability, moral hazard, adverse selection, and makes very few assumptions about the underlying process for returns and information. The paper develops a robust approach to dynamic

^{*}Contact: chassang@princeton.edu.

[†]I'm indebted to Abhijit Banerjee, Roland Benabou, Faruk Gul, Stephen Morris, Wolfgang Pesendorfer, David Sraer and Satoru Takahashi for time, encouragement, and advice. The paper greatly benefited from conversations with Manuel Amador, Bruno Biais, Markus Brunnermeier, Mike Golosov, Johannes Hörner, Augustin Landier, Thomas Mariotti, Ulrich Müller, Guillaume Plantin, Esteban Rossi-Hansberg, Adam Szeidl, Juuso Toikka, Wei Xiong, Muhamet Yildiz and John Zhu. I'm grateful to seminar audiences at Berkeley, the Institute for Advanced Study, Harvard, MIT, NYU, NYU Stern, Paris School of Economics, Princeton, Toulouse School of Economics, Sciences Po Paris, UCL, Warwick, and X-HEC for comments and suggestions. SeHyoun Ahn and Juan Ortner provided excellent research assistance.

¹Throughout the paper, the principal is referred to as she, while the agent is referred to as he.

contracting whose main steps are as follows: 1) identify a simple class of *high-liability* static linear contracts that satisfy attractive and robust efficiency properties; 2) construct *limited-liability* dynamic contracts that achieve the same performance by calibrating the rewards to the agent so that they approximately satisfy key properties of the benchmark high-liability contracts. The resulting dynamic contracts—referred to as calibrated contracts—satisfy attractive performance bounds independently of the underlying process for returns. In particular, the results do not rely on any ergodicity or stationarity assumptions.

The model considers a risk-neutral principal and a risk-neutral agent. Both the principal and the agent are patient. The principal is infinitely lived, while the agent has a large but finite horizon which need not be known to the principal. In every period a fixed amount of resources is to be invested on behalf of the principal by the agent. The agent has private information about the process for returns and can exert costly effort to obtain additional information. The main constraint on contracts is limited liability: the agent cannot receive negative transfers, and rewards are bounded above by per-period resources, which in this framework rules out large deferred payments. The paper makes few assumptions on the underlying probability space and the agent may start with arbitrary private information. Furthermore, the process for information and returns need not be i.i.d. or even ergodic: it may be that with non-vanishing probability there is a large number of periods where returns happen to be negative, or where costly information turns out to be useless.

This is a difficult environment to contract in. The principal is facing both adverse selection (the agent may have persistent private information about returns, or about the cost-effectiveness of information acquisition) and moral hazard (the agent expends effort to acquire information and makes asset allocation decisions). At this level of generality, characterizing optimal contracts is unlikely to be informative and may not actually be possible if the principal has poorly specified beliefs over the environment. Instead the paper develops a robust approach to dynamic contracting which emphasizes prior-free performance bounds.

The first step of the approach relaxes limited liability constraints and identifies a suit-

able high-liability benchmark contract. This benchmark takes the form of a simple linear contract which rewards the agent a share of his externality on the principal. This contract exhibits high-liability since the agent is expected to provide compensation for the losses he causes. While not optimal in every environment, this linear contract constitutes an attractive benchmark: it is max min optimal in an appropriate sense, weakly renegotiation proof, and can guarantee the principal a minimum share of first best returns.

The second step of the approach is to develop a simple class of dynamic contracts that robustly approximate the performance of linear high-liability contracts while satisfying severe limited liability constraints. The key insight is to calibrate both the rewards to the agent and the share of total wealth he is investing so that for all possible strategies and all realizations of uncertainty the payoffs obtained by the agent and the excess returns obtained by the principal remain as tightly linked as they are under benchmark linear contracts. From the perspective of any history, these calibrated contracts induce performance approximately equal to that achieved by the benchmark linear contracts.

Several extensions emphasize the broader applicability of the approach: the methods developed in the paper can be used to approximate a large class of high-liability contracts in addition to the linear benchmark; penalized calibrated contracts can induce uninformed agents to self-screen; multi-agent contracts can be used to jointly manage several agents with different and varying abilities.

The paper hopes to usefully complement the rich literature on optimal dynamic contracting (see for instance Rogerson (1985), Green (1987), Holmström and Milgrom (1987), Spear and Srivastava (1987), Laffont and Tirole (1988), and more recently Battaglini (2005), DeMarzo and Sannikov (2006), Biais et al. (2007, 2010), DeMarzo and Fishman (2007), Sannikov (2008), Edmans et al. (forthcoming) or Zhu (2010)). Because optimal contracts depend finely on the details of the underlying environment, this literature has delivered rich positive predictions on how contract form should vary with the circumstances. However, a limitation of the optimal contracting approach is that it provides little guidance on how well

those contracts perform if the environment is misspecified. The current paper gives up on optimality and develops a class of detail-free contracts that satisfy attractive efficiency properties for a very broad class of stochastic environments. Notably, the performance bounds satisfied by these robust contracts hold in environments where solving for optimal contracts has proved particularly difficult. This includes non-stationary environments (as in Battaglini (2005), Tchisty (2006), He (2009), Pavan et al. (2010) or Garrett and Pavan (2010)), and settings with both moral hazard and adverse selection (such as Sannikov (2007) or Fong (2008)). Still, the contracts developed in the current paper are substantially connected to the optimal contracts derived by DeMarzo and Sannikov (2006) or Biais et al. (2007, 2010) in specific settings. The similarities as well as the differences are instructive and will be discussed in detail.

The paper contributes to a small literature on dynamic design mechanism design with patient players. It is most closely related to the work of Rubinstein (1979), Rubinstein and Yaari (1983) and Radner (1981, 1985) which proves the existence of approximately first-best contracts in a dynamic moral hazard problem where the agent's production function is stationary and common knowledge. More recently Jackson and Sonnenschein (2007) propose simple quota mechanisms that approximately implement any Pareto efficient allocation rule in a class of dynamic multi-agent allocation problems where the agents have i.i.d. preferences. Escobar and Toikka (2009) extend the results of Jackson and Sonnenschein (2007) to the case where preferences follow an irreducible Markov chain. As in these previous approaches, the main idea of the current paper is to constrain payoffs to satisfy key properties that would hold under an ideal benchmark. The central difference between the current paper and Radner (1981, 1985), Jackson and Sonnenschein (2007) or Escobar and Toikka (2009), is that they assume the state of the world follows an ergodic process and their analyses rely strongly on this assumption: the basic idea is to make sure that the empirical distribution of realized outcomes matches the anticipated distribution of outcomes under first-best behavior. This approach is not applicable in the current paper since the underlying environment need not

be ergodic and no law of large numbers need apply.

The methods used in this paper, as well as the emphasis on general stochastic processes, connect the paper to the literature on testing experts (see for instance Foster and Vohra (1998), Fudenberg and Levine (1999), Lehrer (2001) or more recently Al-Najjar and Weinstein (2008), Feinberg and Stewart (2008) and Olszewski and Sandroni (2008)). However, the main question here is not whether good tests are available. Rather, this paper takes a principal-agent approach related to that of Echenique and Shmaya (2007), Olszewski and Peski (2011) or Gradwohl and Salant (2011). These papers show that in such environments there are satisfactory ways to identify experts that generate positive surplus. Olszewski and Peski (2011) relies on ex post high-liability contracts to incentivize truth telling. Gradwohl and Salant (2011) show it is possible to rely on upfront payments instead. Neither paper tackles incentive provision when information acquisition is costly.

The paper is also related to a recent finance literature on appropriate performance measures for wealth managers. Lo (2001), Goetzmann et al. (2007) and Foster and Young (2010) all emphasize the fragility of many performance measures to gaming by the agent, as well as the difficulty of both rewarding and screening agents. In particular Foster and Young (2010) describe environments in which rewarding and screening is in fact impossible. This occurs because their environment allows for a strong form of private saving such that informed managers value income in early periods much more than in later periods.² As a result, talented managers are unwilling to pay the monetary cost needed to induce screening. In contrast, the current paper essentially rules out private savings and considers patient players with constant marginal utility for income. In that case, self-screening can be obtained, even under severe limited liability constraints. The current paper is also related to recent work on the incentive properties of high-watermark contracts. Using Goetzmann et al. (2003)'s valuation of highwatermark contracts as ongoing options, Panageas and Westerfield (2009) show in a specific environment that high-watermark contracts do not necessarily lead to ex-

²Specifically, consumption can be arbitrarily delayed and managers can save on their own at the same rate of returns they generate for the firm.

cessive risk taking when agents have an infinite horizon, although their payoffs are convex in returns. The current paper shows that in fact, a variation on high-watermark contracts can approximately align the interests of the agent and the principal for a large class of underlying stochastic environments.

Finally the paper is related to the literature on robust mechanism design that operationalizes the doctrine set by Wilson (1987), and attempts to characterize mechanisms that behave well under weak assumptions over payoff distributions and beliefs. A rich strand of that literature studies mechanisms that are robust with respect to the solution concept used to characterize the players' behavior.³ The paper is especially related to another strand in this literature, dating back to Hurwicz and Shapiro (1978) and more recently illustrated by Neeman (2003) or Hartline and Roughgarden (2008), which looks for mechanisms that satisfy robust performance bounds over broad sets of fundamentals.⁴ A tricky step, common to this literature and the current paper, is to define appropriate benchmark performance measures that allow for informative worst-case analysis of mechanisms.

The paper is structured as follows. Section 2 describes the framework. Section 3 introduces a benchmark class of high liability linear contracts that satisfy a number of attractive efficiency properties but require high liability. Section 4 is the core of the paper: it develops the idea of calibrated contracts and analyzes their performance. Section 5 generalizes the approach in several ways: first by identifying a broader class of high-liability contracts that can be successfully calibrated; second by providing suitable extensions of the approach when the principal faces multiple potential agents. Section 6 relates calibrated contracts to other contracts of interest and concludes. Appendix A extends the analysis to various environments. Proofs are given in Appendix B, unless mentioned otherwise.

³See for instance Dasgupta et al. (1979), Hagerty and Rogerson (1987), Eliaz (2002), Chung and Ely (2003, 2007) or Bergemann and Morris (2005).

⁴Local approaches are possible and informative. For instance Madarász and Prat (2010) consider screening mechanisms that satisfy strong efficiency bounds for all type distributions within a small neighborhood. Global incentive compatibility constraints play an important role in their analysis, and will also show up in this paper.

2 The Framework

Players, Actions and Payoffs. A principal hires an agent to make investment allocations on her behalf. The agent is active for a large but finite number of periods N . The principal has an infinite horizon and need not know the agent's horizon N . Both the principal and the agent are patient and do not discount future payoffs.⁵

In each period $t \in \{1, \dots, N\}$, the principal invests an amount w at the beginning of the period. The amount of wealth w invested in each period is constant, and can be thought of as a steady state amount of wealth to be invested. The realized wealth w_t after investment is consumed at the end of the period, which rules out private saving. Both the principal and the agent are risk neutral. The agent's outside option is set to zero.

Wealth can be invested in one of K assets whose returns at time t are denoted by $\mathbf{r}_t = (r_{k,t})_{k \in \{1, \dots, K\}}$. Let R denote the set of possible returns \mathbf{r}_t . An asset allocation at time t is a vector $a_t \in A \subset \mathbb{R}^K$ such that $\sum_{k=1}^K a_t = 1$. Set A is convex and compact. It represents constraints on possible allocations. These constraints can be thought of as a mandate set by the principal as in He and Xiong (2010). Let $\langle \cdot, \cdot \rangle$ denote the usual dot product. Given asset allocation a_t and returns \mathbf{r}_t , the consumer's wealth at the end of period t is

$$w_t = w \times (1 + \langle a_t, \mathbf{r}_t \rangle).$$

By assumption, returns are bounded below by -1 so that $w_t \geq 0$ (there cannot be negative resources at the end of the period).

For any pair of allocations $(a, a') \in A^2$, the distance between a and a' is defined by

$$d(a, a') \equiv \sup_{\mathbf{r}_t \in R} | \langle a - a', \mathbf{r}_t \rangle |. \tag{1}$$

The following assumption puts constraints on the set of permissible allocations A and is

⁵Appendix A shows how to extend the analysis when future payoffs are discounted.

maintained throughout the paper.

Assumption 1. *There exists $\bar{d} \in \mathbb{R}^+$ such that for all $(a, a') \in A^2$, $d(a, a') \leq \bar{d}$.*

This assumption limits the magnitude of changes that can occur with the principal getting no feedback.

At the beginning of every period t , managers can expend cost $c_t \in [0, +\infty)$ towards acquiring information. This cost can be the actual cost of obtaining data, an effort cost, or the opportunity cost of time. Managers then make an asset allocation suggestion $a_t \in A$ and receive a payment π_t depending on the realized public history at the end of period t . The manager's objective is to maximize his expected average payoffs

$$\mathbb{E} \left(\frac{1}{N} \sum_{t=1}^N \pi_t - c_t \right). \quad (2)$$

Information. Information acquired at time $t \in \{1, \dots, N\}$ is represented as a random variable I_t from a measurable state space (Ω, σ) to a measurable signal space $(\mathcal{I}, \sigma_{\mathcal{I}})$. Publicly available information is denoted by I_t^0 , includes realized past returns $(\mathbf{r}_s)_{s < t}$, and corresponds to the information available to the principal. In each period, the agent can choose to acquire additional signals $I_t(c)$ at cost $c \in [0, +\infty)$ from a set of possible signals $\{I_t(c) | c \in [0, +\infty)\}$ indexed by their cost. By assumption I_t^0 is measurable with respect to $I_t(0)$ so that the agent is more informed than the principal, regardless of the information he acquires. Given an information acquisition strategy $(c_t)_{t \geq 1}$, let $(\mathcal{F}_t)_{t \geq 1}$ be the manager's filtration (generated by $(I_t(c_t))_{t \geq 1}$), and let $(\mathcal{F}_t^0)_{t \geq 1}$ denote the public information filtration (generated by $(I_t^0)_{t \geq 1}$). The framework allows for adverse selection and moral hazard. At the time of contracting, the agent may already know much more about the process for returns than the principal (through signal $I_0(0)$). Furthermore, the agent's information and information acquisition strategy are private.

For simplicity it is convenient to assume that the principal and the agent have a common

prior P over the state space (Ω, σ) .⁶ Let $\mathcal{P} = (\Omega, \sigma, P)$ denote the resulting probability space (\mathcal{P} will often be referred to as the environment). The paper *does not* assume that either information or returns follow an i.i.d. or ergodic process. This results in a very flexible model. For instance, there may be non-vanishing probability that returns are below their period $t = 1$ expectation for an arbitrarily large number of periods. Also, the value of the information that managers can collect may vary in arbitrary ways. For instance, once valuable trading strategies can become obsolete over time.

Strategies. Altogether, an agent's strategy consists of an information acquisition strategy $c = (c_t)_{t \in \mathbb{N}}$, and an asset allocation strategy $a = (a_t)_{t \in \mathbb{N}}$, where both c_t and a_t are adapted to the information available to the manager at the time of decision. Let a_t^0 and a_t^* respectively denote efficient asset allocations under information \mathcal{F}_t^0 and \mathcal{F}_t :

$$a_t^0 \in \arg \max_{a \in A} \mathbb{E}[\langle a, \mathbf{r}_t \rangle | \mathcal{F}_t^0] \quad \text{and} \quad a_t^* \in \arg \max_{a \in A} \mathbb{E}[\langle a, \mathbf{r}_t \rangle | \mathcal{F}_t]. \quad (3)$$

Allocation a_t^0 is the allocation the principal could pick on her own, given public information \mathcal{F}_t^0 . Let $w_t^0 = w \times (1 + \langle a_t^0, \mathbf{r}_t \rangle)$ and $w_t = w \times (1 + \langle a_t, \mathbf{r}_t \rangle)$ denote realized wealth under allocation a_t^0 and under the allocation a_t actually chosen by the agent.

Contracts. Contracts $(\pi_t)_{t \in \mathbb{N}}$ are adapted to public histories observed by the principal, where public histories consist of past public information (including past returns) as well as past suggested asset allocations by the agent. The principal has commitment power but transfers are subject to resource constraints: in every period t ,

$$0 \leq \pi_t \leq w_t, \quad (4)$$

The constraint that $0 \leq \pi_t$ corresponds to a limited-liability constraint on the agent's side: the agent does not have access to side resources that can be pledged in the contract. Sym-

⁶Results extend to a non-common prior setting, taking expectations under the agent's prior.

metrically, the constraint that $\pi_t \leq w_t$ limits payments to the agent to resources generated in that period. This limits how long the payment of wages can be delayed and precludes the possibility of large deferred payments.

These constraints are at the origin of the contracting problem: the agent does not share on the downside, and rewards must be given in real time rather than delayed until the end. Clearly, these are strong liability constraints which may be relaxed in a number of settings. Naturally, the efficiency bounds derived in the paper remain valid when limited liability constraint (4) is relaxed.

3 A High-Liability Benchmark

The environment described in Section 2 involves both moral hazard and adverse selection: the agent must acquire information and makes asset allocation decisions that may or may not benefit the principal; in addition the information that the agent has or may acquire is private. At this level of generality, informative characterizations of optimal dynamic contracts are unlikely and solving for optimal contracts may also be of limited use if the principal doesn't have well defined beliefs over the underlying environment.

The paper embraces an alternative approach to dynamic contracting which aims to identify contracts satisfying robust efficiency properties over broad classes of environments. The first step of the analysis defines a class of benchmark contracts that have attractive efficiency properties, but violate limited liability constraint (4). The second step of the analysis constructs a class of dynamic contracts that satisfy constraint (4), and achieve performance approximately as good as that of the benchmark contracts, *regardless of the underlying environment* \mathcal{P} .

Benchmark contracts. The contracts used as benchmark are linear contracts in which the agent's reward π_t in period t is a share α of the externality his decisions have on the

principal:

$$\forall t, \quad \pi_t = \alpha(w_t - w_t^0).^7 \tag{5}$$

These benchmark contracts are attractive for the following reasons:

- (i) they satisfy a robust efficiency bound regardless of the underlying environment and are max min optimal in an appropriate sense;
- (ii) they are weakly renegotiation proof in the sense of Bernheim and Ray (1989) and Farrell and Maskin (1989);
- (iii) they are the only class of contracts satisfying a demanding no-loss condition, under which the agent makes positive profits if and only if the principal gets positive surplus

Note that even though both parties are risk-neutral, the fact that the agent has significant private information means that fixed-price contracts in which all productive assets are sold to the agent need not be optimal.⁸

Robust efficiency properties. Recent work by Rogerson (2003), Chu and Sappington (2007) and Bose et al. (2011) has emphasized that simple contracts can often guarantee large shares of the second-best surplus in contexts ranging from procurement to principal-agent problems. These approaches have focused on parametric models for which it is possible to compute the second-best explicitly.

Theorem 1, stated below, contributes to this literature by providing a non-parametric bound on the share of first-best surplus that can be obtained through linear contracts, and showing that benchmark contracts are max min optimal with respect to an appropriate class of environments.

⁷Recall that w_t and w_t^0 respectively denote final wealth under the agent's suggested asset allocation and under the default, public information, asset allocation. For instance, if $\alpha = 20\%$ and the default allocation a_t^0 is to invest all wealth in risk-free bonds, the benchmark contract pays the agent 20% of the excess-returns when he beats the risk-free rate, and charges him 20% of the foregone returns when he under-performs the risk-free rate.

⁸In the max min problem defined below, fixed price contracts do not robustly guarantee a positive share of first-best surplus: any positive price will cause the agent not to participate in some environment.

To show this, additional notation must be introduced. Given a contract $\pi = (\pi_t)_{t \geq 0}$, the agent solves optimization problem

$$\max_{c,a} \mathbb{E} \left(\frac{1}{N} \sum_{t=1}^N \pi_t - c_t \right). \quad (\text{P1})$$

The corresponding per-period excess returns r_π accruing to the principal (net of payments to the agent) are

$$r_\pi \equiv \inf \left\{ \mathbb{E}_{c,a} \left(\frac{1}{Nw} \sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \middle| (c,a) \text{ solves (P1)} \right\}.$$

The expression for returns r_π involves an inf since the agent may be indifferent between multiple strategy profiles. Returns accruing to the principal when the contract is $\pi_t = \alpha(w_t - w_t^0)$ are denoted by r_α . In anticipation of technical subtleties to come, it is useful to note that because the underlying environment is very general, the paper cannot rule out binding global incentive compatibility constraints.

For any $\hat{c} \in [0, +\infty)$, let $r_{\max}(\hat{c})$ denote the production function for returns, i.e. expected per-period returns generated when the agent: 1) incurs an expected per-period cost of information acquisition equal to \hat{c} ; 2) chooses optimal asset allocation a^* given information; and 3) requires no rewards. Formally we have

$$r_{\max}(\hat{c}) \equiv \sup_{\substack{c \text{ s.t.} \\ \mathbb{E}[\frac{1}{N} \sum_{t=1}^N c_t] \leq \hat{c}}} \mathbb{E}_{c,a^*} \left(\frac{1}{N} \sum_{t=1}^N \langle a_t^* - a_t^0, \mathbf{r}_t \rangle \right).$$

First best surplus corresponds to maximizing $w r_{\max}(\hat{c}) - \hat{c}$. Note that since r_{\max} is bounded above by \bar{d} , the optimum is necessarily attained for a finite value of \hat{c} . Denote by r_{FB} and c_{FB} the first-best average per-period returns and first-best average per-period costs. For any $\rho \leq 1$, denote by

$$\mathbb{P}_\rho = \left\{ \mathcal{P} \mid \frac{c_{FB}}{w r_{FB}} \leq \rho \right\}$$

the set of environments such that the ratio of costs to returns at first-best is bounded above by ρ . Note that by definition, it must be that $\frac{c_{FB}}{wr_{FB}} \leq 1$.

Theorem 1 (robust efficiency bounds). *(i) For any probability space \mathcal{P} ,*

$$wr_\alpha \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(wr_{\max}(\hat{c}) - \frac{c}{\alpha} \right). \quad (6)$$

(ii) For any $\rho \in (0, 1)$, the benchmark contract of parameter $\alpha = \sqrt{\rho}$ satisfies

$$\begin{aligned} \max_{\pi = (\pi_t)_{t \geq 1}} \min_{\mathcal{P} \in \mathbb{P}_\rho} \frac{wr_\pi}{wr_{FB} - c_{FB}} &= \min_{\mathcal{P} \in \mathbb{P}_\rho} \frac{wr_\alpha}{wr_{FB} - c_{FB}} \\ &= 1 - 2 \frac{\sqrt{\rho}}{1 + \sqrt{\rho}}. \end{aligned} \quad (7)$$

Given a benchmark contract $\pi_t = \alpha(w_t - w_t^0)$, point (i) provides a lower bound for the returns that the principal obtains for any environment \mathcal{P} . Point (ii) leverages point (i) to show that the linear contract of parameter $\alpha = \sqrt{\rho}$ guarantees the highest possible proportion of first-best surplus over environments $\mathcal{P} \in \mathbb{P}_\rho$. Specifically, the benchmark contract guarantees a proportion $1 - 2 \frac{\sqrt{\rho}}{1 + \sqrt{\rho}}$ of first-best surplus.⁹

Note that although benchmark contracts are static, i.e. they depend only on current outcomes, Theorem 1 shows that benchmark contracts are max min optimal among all possible dynamic contracts. This provides foundations for the use of linear contracts as a benchmark. If one is willing to further restrict the set of possible probability spaces \mathcal{P} against which contract performance is evaluated, other contracts may be more attractive. Section 5 returns briefly to this point.

Other remarkable properties. As was noted previously, benchmark contracts have additional merits. The following results hold.

⁹Without scaling by first-best returns, the maxmin problem described by (7) would be degenerate: because first-best returns can be arbitrarily low, unscaled maxmin returns are equal to 0.

Fact 1. *Benchmark contracts are weakly renegotiation proof in the sense of Bernheim and Ray (1989) and Farrell and Maskin (1989).*

Indeed, since benchmark contracts are independent of history, the principal and the agent are never tempted to renegotiate to a continuation contract starting from a different history.

Finally, benchmark contracts satisfy the following no-loss property: for any strategy (c, a) such that the agent obtains positive profit – even suboptimal ones – the principal must also obtain positive surplus.¹⁰

Fact 2 (no loss). *Under the benchmark contract, for all environment \mathcal{P} and all strategy profiles (c, a) , $\mathbb{E}_{c,a} \left[\sum_{t=1}^N \pi_t \right] \geq 0 \iff \mathbb{E}_{c,a} \left[\sum_{t=1}^N w_t - w_t^0 - \pi_t \right] \geq 0.$*

The converse holds. If a contract $(\pi_t)_{t \geq 1}$ is such that for all \mathcal{P} and all strategy profiles (c, a) , $\mathbb{E}_{c,a} \left[\sum_{t=1}^N \pi_t \right] \geq 0 \iff \mathbb{E}_{c,a} \left[\sum_{t=1}^N w_t - w_t^0 - \pi_t \right] \geq 0$, then there exists $\alpha \in (0, 1)$ such that for all t , $\pi_t = \alpha(w_t - w_t^0).$

This property is attractive especially in environments where the agent may not be fully optimizing: provided that a suboptimal strategy profile does not generate negative profit for the agent, it can only benefit the principal. While the bulk of the paper assumes that agents are rational and fully optimizing, Appendix A returns to the question of contract performance when agents can be temporarily irrational.

Theorem 1 and Facts 1 and 2 motivate the use of linear contracts as a robust benchmark. Unfortunately benchmark contracts do not satisfy limited liability constraint (4). The next section constructs equally robust dynamic contracts that perform approximately as well as benchmark contracts in all these respects, while also satisfying limited liability constraints (4).

¹⁰This is a form of robustness to the solution concept which allows for “faulty” non-best-replying agents, as in Eliaz (2002).

4 Calibrated Contracts

This section introduces a novel class of dynamic limited-liability contracts, referred to as calibrated contracts, which robustly approximate the performance of high-liability benchmark contracts. Calibrated contracts are analyzed in two steps. Section 4.1 describes the contract and states the main efficiency result. Section 4.2 details the mechanics underlying calibrated contracts.

4.1 The Contract

In every period t , the agent is allowed to invest a share $\lambda_t \in [0, 1]$ of the principal's wealth, while the remaining share $1 - \lambda_t$ is invested in the default asset allocation a_t^0 . At the end of each period t , the agent receives a payment π_t .¹¹

Specifying investment shares and rewards $(\lambda_t, \pi_t)_{t \geq 1}$ requires additional notation. For all periods T and $T' < T$, define

$$\Pi_T = \sum_{t=1}^T \pi_t; \quad \Sigma_T = \sum_{t=1}^T w_t - w_t^0; \quad S_T = \sum_{t=1}^T \lambda_t (w_t - w_t^0) \quad (8)$$

and

$$\Sigma_{T \setminus T'} = \sum_{t=T'}^T w_t - w_t^0; \quad S_{T \setminus T'} = \sum_{t=T'}^T \lambda_t (w_t - w_t^0). \quad (9)$$

Value Π_T corresponds to the payoffs that the agent has obtained; Σ_T corresponds to the excess returns that would have been generated by fully investing according to the agent's suggested asset allocation; S_T corresponds to the actual excess returns that have been generated given that only a share λ_t of wealth w is invested according to the agent's suggestion. Values $\Sigma_{T \setminus T'}$ and $S_{T \setminus T'}$ compute the same quantities over time range $\{T', \dots, T\}$.

The difference $\Sigma_{T \setminus T'} - S_{T \setminus T'} = \sum_{t=T'}^T (1 - \lambda_t)(w_t - w_t^0)$ corresponds to the foregone gains from investing only a share λ_t of resources according to the agent's allocation between T'

¹¹Note that under the benchmark high-liability contract, in equilibrium, it must be that for all $t \geq 1$, $\mathbb{E}(w_t - w_t^0) \geq 0$. This is why under the benchmark contract, parameter λ_t is optimally set to 1.

and T . The difference $\Pi_T - \alpha S_T$ corresponds to the agent's excess rewards, the target being to reward him a share α of his externality S_T on the principal. We can now define calibrated contracts formally.

Using the notation $(x)^+ = \max\{0, x\}$, investment shares and rewards $(\lambda_t, \pi_t)_{t \geq 1}$ are defined recursively as follows. Initially set $\lambda_1 = 1$, $\pi_1 = 0$.¹² For all subsequent periods $T \geq 1$, let

$$\begin{aligned} \lambda_{T+1} &\equiv \frac{\alpha [\max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}]^+}{[\Pi_T - \alpha S_T]^+ + \alpha [\max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}]^+} \\ &\equiv \frac{\alpha \times \text{maximum foregone gain}}{\text{excess rewards} + \alpha \times \text{maximum foregone gain}} \end{aligned} \quad (10)$$

with the convention that $\frac{0}{0} = 1$, and

$$\begin{aligned} \pi_{T+1} &\equiv \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if } \Pi_T \leq \alpha S_T \\ 0 & \text{otherwise} \end{cases} \\ &\equiv \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if } \text{rewards} \leq \alpha \times \text{actual excess returns} \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (11)$$

Note that the contract specified above satisfies limited liability constraint (4): payments π_t are positive and bounded above by w_t . Theorem 2 (stated below) shows that as horizon N grows large, this class of contracts approximates the performance of benchmark high liability contracts. Some additional notation is needed. Given a contract specification $(\lambda, \pi) = (\lambda_t, \pi_t)_{t \geq 1}$, let $r_{\lambda, \pi}$ denote the net excess returns delivered by the agent under contract (λ, π) :

$$r_{\lambda, \pi} = \inf \left\{ \mathbb{E}_{c, a} \left(\frac{1}{Nw} \sum_{t=1}^N \lambda_t (w_t - w_t^0) - \pi_t \right) \middle| (c, a) \text{ solves } \max_{c, a} \mathbb{E} \left(\sum_{t=1}^N \pi_t - c_t \right) \right\}.$$

¹²The choice of initial conditions does not affect long term performance.

For any history h_T , normalized net returns conditional on h_T are

$$r_{\lambda,\pi}|h_T = \inf \left\{ \mathbb{E}_{c,a} \left(\frac{1}{Nw} \sum_{t=T+1}^N \lambda_t(w_t - w_t^0) - \pi_t \middle| h_T \right) \middle| (c, a) \text{ solves } \max_{c,a} \mathbb{E} \left(\sum_{t=1}^N \pi_t - c_t \right) \right\}.$$

When the contract in question is the benchmark linear contract of parameter α , net returns accruing to the principal are denoted by r_α (similarly, let $r_\alpha|h_T$ denote conditional returns at history h_T).

Theorem 2 (approximate performance). *Pick $\alpha_0 \in (0, 1)$ and for any $\eta \in (0, 1)$, let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. Consider the calibrated contract (λ, π) defined by (10) and (11). There exists a constant m independent of time horizon N and probability space \mathcal{P} such that,*

$$r_{\lambda,\pi} \geq (1 - \eta)r_{\alpha_0} - m \frac{1}{\sqrt{N}} \tag{12}$$

$$\forall h_T, \quad r_{\lambda,\pi}|h_T \geq (1 - \eta)r_{\alpha_0}|h_T - m \frac{1}{\sqrt{N}}. \tag{13}$$

It follows that for N large enough, the calibrated contract described by (10) and (11) generates a share approximately $1 - \eta$ of the returns the principal obtains under the benchmark contract of parameter α_0 . Appendix A provides a similar, though weaker, bound when future payoffs are discounted. Note that this result involves two approximations. First there is a multiplicative loss $1 - \eta$ which comes from sharing an additional proportion η of profits with the agent. Second there is an additive loss of order $1/\sqrt{N}$ which comes from imperfectly approximating incentives. The mechanics underlying Theorem 2, and the reason why an additional incentive η is needed will be discussed in details in Section 4.2.

Note that for Theorem 2 to hold, it is not sufficient to just reward the agent according to the payment rule $(\pi_t)_{t \geq 1}$ defined by (10) and (11). It is important that the principal actually invest shares $(\lambda_t)_{t \geq 1}$ of her wealth according to the agent's suggestion. Indeed the reward scheme $(\pi_t)_{t \geq 1}$ does not to induce perfectly good behavior from the agent.¹³ Rather, the

¹³For instance, an agent who has lost or never had any informational advantage may pick allocations a_t

payment scheme $(\pi_t)_{t \geq 1}$ reduces misbehavior to the point where it can be resolved by using the cautious investment rule specified by $(\lambda_t)_{t \geq 1}$.

Finally, note that returns $r_{\lambda, \pi}$ and r_{α_0} are computed under the assumption that the agent's behavior is an exact best-reply. No approximate best-reply assumption is made. Appendix A shows that calibrated contracts continue to perform well if rationality is weakened so that the agent may behave suboptimally over an arbitrary interval of time.

4.2 The Mechanics of Calibrated Contracts

This idea behind calibrated contracts is to identify key incentive properties that hold under the benchmark contract and calibrate payments $(\pi_t)_{t \geq 1}$ to the agent as well as investment shares $(\lambda_t)_{t \geq 1}$ so that the same incentive properties are approximately satisfied under the calibrated contract. The properties of benchmark contracts that calibrated contracts attempt to satisfy are as follows. For all histories h_T , benchmark contracts are such that

$$\Pi_T = \alpha S_T \quad (\text{no excess rewards}) \quad (14)$$

$$\forall T' \leq T, \quad \Sigma_{T \setminus T'} \leq S_{T \setminus T'} \quad (\text{no foregone performance}). \quad (15)$$

In words, the agent receives a share α of his actual performance S_T , and over any time interval $\{T', \dots, T\}$, his actual performance $S_{T \setminus T'}$ (although potentially hindered by investment shares $\lambda_t \leq 1$) is at least as high as his potential performance $\Sigma_{T \setminus T'}$.¹⁴ Note that for any T , the family of inequalities (15) can be summarized by the single inequality

$$\max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'} \leq 0.$$

that are inferior to a_t^0 , simply because they are different and, through volatility, induce a non-zero probability of reward. The investment rule $(\lambda_t)_{t \geq 1}$ insulates the principal from such misbehavior.

¹⁴To obtain only inequality (12) in Theorem 2, it would be sufficient to consider only inequality $\Sigma_T \leq S_T$ rather than the full family of inequalities described by (15). Considering the full family of inequalities (15) yields the history-independent performance bound (13).

Let us now show how (10) and (11) calibrate parameters $(\lambda_t, \pi_t)_{t \geq 1}$ so that these properties hold approximately, while satisfying limited liability constraint (4). Define regrets

$$\begin{aligned}\mathcal{R}_{1,T} &\equiv \Pi_T - \alpha S_T \\ \mathcal{R}_{2,T} &\equiv \max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}.\end{aligned}$$

Regret $\mathcal{R}_{1,T}$ measures how overpaid the agent has been, while regret $\mathcal{R}_{2,T}$ measures maximum foregone profits from not fully investing according to the agent’s allocation. The goals are: 1) to keep $\mathcal{R}_{1,T}$ small so that the agent’s reward Π_T is a share approximately α of his actual externality S_T on the principal; 2) to keep $\mathcal{R}_{2,T}^+$ small, so that the foregone returns are not large.

These goals can be achieved by following the regret-minimization approach of Blackwell (1956) and Hannan (1957).¹⁵ Define $\mathcal{R}_T \equiv (\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ and $\rho_T \equiv \mathcal{R}_T - \mathcal{R}_{T-1}$ the vector of flow regrets.¹⁶ In order to keep regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ small, it is sufficient to keep vector \mathcal{R}_T small. This can be achieved by choosing sequences $(\pi_t)_{t \geq 1}$ and $(\lambda_t)_{t \geq 1}$ so that

$$\forall T \geq 1, \forall w_{T+1}, \forall w_{T+1}^0, \quad \langle \rho_{T+1}, \mathcal{R}_T^+ \rangle \leq 0. \quad (16)$$

Inequality (16) is known as an approachability condition, and ensures that flow regrets ρ_{T+1} point in the direction opposite to that of aggregate regrets \mathcal{R}_T . This puts strong bounds on the speed at which aggregate regrets $(\mathcal{R}_T)_{T \geq 1}$ can grow.

By construction, regret $(\mathcal{R}_{2,T})_{T \geq 1}$ —which measures maximum foregone gains—satisfies

$$\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T}^+.¹⁷$$

¹⁵See also Foster and Vohra (1999) or Cesa-Bianchi and Lugosi (2006). Regret measure $\mathcal{R}_{2,T}$ is related to “tracking” regrets (Cesa-Bianchi and Lugosi, 2006).

¹⁶Vector \mathcal{R}_T is defined as $(\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ rather than $(\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ only because it leads to a slight improvement in performance bounds.

¹⁷This implies that regret $\mathcal{R}_{2,T}$ can be computed using at most $O(T)$ operations.

Thus, condition (16) is equivalent to

$$[\pi_{T+1} - \alpha\lambda_{T+1}(w_{T+1} - w_{T+1}^0)]\mathcal{R}_{1,T}^+ + \alpha^2 [(1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0) + \mathcal{R}_{2,T}^+ - \mathcal{R}_{2,T}] \mathcal{R}_{2,T}^+ \leq 0.$$

Using the identity $(\mathcal{R}_{2,T}^+ - \mathcal{R}_{2,T})\mathcal{R}_{2,T}^+ = 0$, it follows that approachability condition (16) is equivalent to

$$\begin{aligned} & [\pi_{T+1} - \alpha\lambda_{T+1}(w_{T+1} - w_{T+1}^0)]\mathcal{R}_{1,T}^+ + \alpha^2 [(1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0)] \mathcal{R}_{2,T}^+ \leq 0 \\ \iff & \pi_{T+1}\mathcal{R}_{1,T}^+ - [\alpha\lambda_{T+1}\mathcal{R}_{1,T}^+ - \alpha^2(1 - \lambda_{T+1})\mathcal{R}_{2,T}^+](w_{T+1} - w_{T+1}^0) \leq 0. \end{aligned}$$

Hence approachability condition (16) can be satisfied for any realization of w_{T+1} and w_{T+1}^0 by setting

$$\lambda_{T+1} = \frac{\alpha\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha\mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} \alpha\lambda_{T+1}(w_{T+1} - w_{T+1}^0)^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{if } \mathcal{R}_{1,T} > 0 \end{cases}$$

which corresponds to the calibrated contract defined by (10) and (11).

The following lemma shows that under the contract defined by (10) and (11), incentive properties (14) and (15) are approximately satisfied.

Lemma 1 (approximate incentives). *For all T , all $T' \leq T$ and all possible histories,*

$$\Sigma_{T \setminus T'} - S_{T \setminus T'} \leq w\bar{d}\sqrt{T} \tag{17}$$

$$-\alpha w\bar{d} \leq \Pi_T - \alpha S_T \leq \alpha w\bar{d}\sqrt{T}. \tag{18}$$

Lemma 1 implies that incentive properties (14) and (15) hold at any possible history h_T , up to an error term of order \sqrt{T} , which is small compared to the number of periods T . Note that this holds sample path by sample path, rather than in expectation or in equilibrium.

Proof. Let $d_t = \sup_{\mathbf{r} \in R} |a_t - a_t^0, \mathbf{r}|$ denote the magnitude of positions taken by the agent

in period t . We first show that $\|\mathcal{R}_T^+\|^2 \leq \alpha^2 w^2 \sum_{t=1}^T d_t^2$. The proof is by induction. The property clearly holds at $T = 1$. Assume it holds at T . Consider the case where $\mathcal{R}_{2,T} > 0$ (i.e. there are some foregone returns). Since approachability condition (16) holds, we have that

$$\begin{aligned} \|\mathcal{R}_{T+1}^+\|^2 &\leq \|\mathcal{R}_T^+ + \rho_{T+1}\|^2 = \|\mathcal{R}_T^+\|^2 + 2\langle \mathcal{R}_T^+, \rho_{T+1} \rangle + \|\rho_{T+1}\|^2 \\ &\leq \|\mathcal{R}_T^+\|^2 + \|\rho_{T+1}\|^2. \end{aligned}$$

In addition $\|\rho_{T+1}\|^2 \leq \alpha^2 \lambda_{T+1}^2 (w_{T+1} - w_{T+1}^0)^2 + \alpha^2 (1 - \lambda_{T+1})^2 (w_{T+1} - w_{T+1}^0)^2 \leq \alpha^2 w^2 d_{T+1}^2$. Altogether this shows that the induction hypothesis holds when $\mathcal{R}_{2,T} > 0$. A similar proof holds when $\mathcal{R}_{2,T} < 0$, taking into account that in this case, $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})(w_{T+1} - w_{T+1}^0)$. Hence, by induction, this implies that for all $T \geq 1$, $\|\mathcal{R}_T^+\|^2 \leq \alpha^2 w^2 \sum_{t=1}^T d_t^2$. This proves (17) and the right-hand side of (18).

The left-hand side of (18) is also proven by induction. If $\Pi_T \in [\alpha S_T - \alpha w \bar{d}, \alpha S_T]$, then $\mathcal{R}_{1,T}^+ = 0$, $\lambda_{T+1} = 1$, and $\pi_{T+1} = \alpha(w_{T+1} - w_{T+1}^0)^+$. Hence by construction, $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha w \bar{d}$. If instead $\Pi_T > \alpha S_T$, then by definition of \bar{d} , $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha w \bar{d}$. This implies the left-hand side of (18). \square

As the next lemma shows, the approximate incentive conditions given by Lemma 1 imply performance bounds for calibrated contracts.

Lemma 2. *Pick $\alpha_0 \in (0, 1)$ and for any $\eta \in (0, 1)$ let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. Consider a contract (λ, π) and numbers A, B and C such that for all final histories h_N , $\Sigma_N - S_N \leq A$ and $-B \leq \Pi_N - \alpha S_N \leq C$. Then*

$$r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - \frac{1}{Nw} \left[C + \frac{1 - \eta}{\eta} (\alpha A + B + C) \right].$$

Theorem 2 is an immediate corollary of Lemmas 2 and 1. Intuitively, Lemma 1 shows that the calibrated contract (λ, π) defined by (10) and (11) gets incentives approximately right.

Lemma 2 implies that when incentives are approximately right, then performance must be approximately right as well. This is not an obvious result since approximation errors with respect to incentives may cause the agent to change his behavior significantly. Indeed, whenever global incentive constraints are binding or almost binding under the benchmark linear contract of parameter α_0 , getting incentives slightly wrong will result in large shifts in behavior and poor performance. For instance, this would be the case if under the benchmark contract, the agent were indifferent between working hard and not working at all. For this reason incentives must be reinforced. By sharing an additional fraction η of her returns, the principal ensures that potential changes in the agent’s behavior do not compromise performance. Madarász and Prat (2010) make the same point in a screening context.

5 Extensions

This section presents three important extensions. Section 5.1 shows how the methods described in Section 4 can be used to calibrate a broader class of high-liability contracts. Section 5.2 shows how to induce uninformed agents to self-screen. Section 5.3 proposes a multi-agent extension of the contracts described in Sections 3 and 4.

5.1 Calibrating a broader class of contracts

Section 3 motivates the use of linear benchmark contracts by showing that they are max min optimal against classes of environments \mathcal{P} such that the ratio of costs to returns at first-best is bounded above. If we are willing to consider smaller or different classes of environments, different benchmarks may be desirable.

Exploring how different classes of environments \mathcal{P} map into different max min optimal high liability contracts is beyond the scope of this paper. Rather, this section takes as given a high-liability contract with aggregate rewards denoted by Π_N^0 , where Π_N^0 is adapted to the principal’s information at time N . The question is whether high-liability contract Π_N^0

can be calibrated using limited-liability contracts. The following assumption turns out to be sufficient.

Assumption 2. Benchmark contract Π_N^0 can be written as $\Pi_N^0 = \sum_{t=1}^N \pi_t^0$, such that

(i) π_t^0 is adapted to the information available to the principal at time t ;

(ii) $w_t = w_t^0$ implies $\pi_t^0 \geq 0$;

(iii) $\pi_t^0 \leq w_t$ and there exists $\bar{\pi} > 0$ such that, $\sup |\pi_t^0| \leq \bar{\pi}$.

It is immediate that Assumption 2 holds for all contracts of the form $\Pi_N^0 = \sum_{t=1}^N \alpha_t^0 (w_t - w_t^0)$ where $\alpha_t^0 \in (0, 1)$ is adapted to public information. Assumption 2 also holds for contracts of the form

$$\Pi_N^0 = G \left(\sum_{t=1}^N \phi(w_t - w_t^0) \right)$$

where $\phi(0) = G(0) = 0$ and G and ϕ are Lipschitz, with constants κ_G and κ_ϕ such that $\kappa_G \kappa_\phi \leq 1$. For instance, this includes contracts such that the agent gets paid a positive reward only when returns are above a threshold, i.e contracts such that

$$\Pi_N^0 = \begin{cases} \alpha \Sigma_N & \text{if } \Sigma_N < 0 \\ 0 & \text{if } \Sigma_N \in [0, \underline{\Sigma}] \\ \alpha(\Sigma_N - \underline{\Sigma}) & \text{if } \Sigma_N > \underline{\Sigma} \end{cases} . \quad (19)$$

Calibration. Theorem 3, stated below, shows that the performance of any contract satisfying Assumption 2 can be robustly approximated using dynamic limited liability contracts.

As in Section 4 an additional incentive wedge is necessary. For any $\eta > 0$ define the auxiliary contract

$$\pi_t^\eta \equiv \pi_t^0 + \eta(w_t - w_t^0 - \pi_t^\eta) = \frac{1}{1+\eta} \pi_t^0 + \frac{\eta}{1+\eta} (w_t - w_t^0).$$

If contract $(\pi_t^0)_{t \geq 1}$ satisfies Assumption 2, then so does contract $(\pi_t^\eta)_{t \geq 1}$. In particular, $|\pi_t^\eta| \leq \frac{1}{1+\eta} \bar{\pi} + \frac{\eta}{1+\eta} w \bar{d} \equiv \bar{\pi}^\eta$.

The approach consists in calibrating the incentives provided by contract $(\pi_t^\eta)_{t \geq 1}$. The two instruments used are rewards $(\pi_t)_{t \geq 1}$ and the proportion of resources $(\lambda_t)_{t \geq 1}$ managed by the agent. Define $\pi_t^\eta(\lambda_t) = \lambda_t \pi_t^\eta$. The regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ to be minimized are:

$$\mathcal{R}_{1,T} = \sum_{t=1}^T \pi_t - \pi_t^\eta(\lambda_t) \quad (\text{no excess rewards}) \quad (20)$$

$$\mathcal{R}_{2,T} = \max_{T' \leq T} \sum_{t=T'}^T \pi_t^\eta - \pi_t^\eta(\lambda_t) \quad (\text{no foregone performance}). \quad (21)$$

The usual approachability condition yields contract parameters $(\lambda_t, \pi_t)_{t \geq 1}$ of the form,

$$\lambda_{T+1} = \frac{\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} [\pi_{T+1}^\eta(\lambda_{T+1})]^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise.} \end{cases} \quad (22)$$

As in Section 4, and as detailed in Appendix B, this ensures that the vector of regrets $(\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ remains of order \sqrt{T} , so that incentives are approximately correct. The following performance bounds obtain.

Theorem 3. *There exists a constant m independent of environment \mathcal{P} and time horizon N , such that under contract $(\lambda_t, \pi_t)_{t \geq 1}$, returns accruing to the principal satisfy*

$$r_{\lambda, \pi} \geq \frac{1}{1 + \eta} r_{\pi^0} - m \frac{1}{\sqrt{N}} \quad (23)$$

$$\forall h_T, \quad r_{\lambda, \pi} | h_T \geq \frac{1}{1 + \eta} r_{\pi^0} | h_T - m \frac{1}{\sqrt{N}} \quad (24)$$

Hence any contract satisfying Assumption 2 can be calibrated using the methods of Section 4. This include contracts such that the agent is paid only after returns reach a minimum threshold, as in (19). What constitutes an appropriate high liability benchmark will depend on the class of environments against which one wishes to measure efficiency. Exploring such classes of environments and the corresponding max min optimal high-liability contracts in left for future research.

5.2 Inducing self-screening by uninformed agents

The framework of Section 2 allows for arbitrary adverse selection. In particular, under some realizations, the agent may learn ex ante that he has no ability to generate valuable information. This section shows how to induce such uninformed agents to self-screen, i.e. to not participate. This is important when most agents are uninformed and only a handful can generate positive value added. The next section shows how to extend the approach of Sections 3 and 4 to jointly manage a small number of agents.

An agent is said to be uninformed if conditional on his initial information $I_0(0)$ he cannot generate positive returns.¹⁸ An issue with the calibrated contract defined by (10) and (11) is that by construction rewards are positive and a sufficiently long-lived uninformed agent can obtain significant expected payoffs from luck and volatility alone.¹⁹ In order to induce entirely uninformed agents to self-screen, i.e. to not participate in the first place, some amount of liability is required. The question is how much? Surprisingly, it turns out that for any liability level b available in the first period, a contract can be found that robustly induces uninformed agents to self-screen at a minimal efficiency cost, regardless of the agents' time horizon, or of environment \mathcal{P} .

Specifically, screening is induced by first imposing an initial participation cost $-b$ on the agent, and then only paying the agent when his performance is above a dynamic hurdle Θ_T which depends on the size of positions he has been taking. Given a free parameter $M > 0$, define

$$\Theta_T \equiv 2w \left(1 + \sqrt{\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2} \right) \sqrt{M + \ln \left(\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2 \right)}, \quad (25)$$

where $d_t = \sup_{\mathbf{r}_t \in R} | \langle a_t - a_t^0, \mathbf{r}_t \rangle |$ and $\lambda_t d_t$ measures the size of the agent's effective bet

¹⁸Formally, $\max_{c,a} \mathbb{E}_{c,a} \left(\sum_{t=1}^N w_t - w_t^0 \middle| I_0(0) \right) = 0$.

¹⁹Even if the agent has no information and all assets have the same expected returns, systematically picking assets different from the benchmark allocation will allow the agent to obtain rewards of order \sqrt{N} with non-vanishing probability.

$\lambda_t(a_t - a_t^0)$ away from the default allocation a_t^0 (note that by Assumption 1, $d_t \leq \bar{d}$). Hurdle Θ_T is an aggregate measure of how active the agent has been. If the agent makes significant bets away from a_t^0 in every period then Θ_T will be of order $\sqrt{T \ln T}$. If the agent makes few bets, hurdle Θ_T will remain small. The quantity $\bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2$ is a measure of time under which $(\Sigma_T)_{T \geq 1}$ will have at most the variation of a standard Brownian motion.

Hurdled calibrated contracts are defined by a sequence $(\lambda_t, \pi_t, \pi_t^\Theta)_{t \geq 1}$. The sequence $(\lambda_t, \pi_t)_{t \geq 1}$ is still defined according to recurrence equations (10) and (11), and λ_t is still the share of wealth actually invested by the agent. However, for $t \geq 1$, reward π_t is no longer paid to the agent for sure. Rather, the agent is paid a hurdled reward π_t^Θ such that $\pi_1^\Theta = -b$, and $\pi_t^\Theta = \mathbf{1}_{S_t \geq \Theta_t} \pi_t$, i.e. potential reward π_t is paid to the agent if and only if the surplus S_t he has generated is greater than hurdle Θ_t .

An intuitive rationale for the form of hurdle Θ_T is as follows. Let the agent be uninformed, so that the process $(S_t)_{t \geq 1}$ is at best a martingale, and imagine that the agent is frequently active, i.e. $\sum_{t=1}^T d_t^2$ is of order T . Then hurdle Θ_T is of order $\sqrt{T \ln T}$. The law of the iterated logarithm implies that with probability 1, as T gets large, $\max_{T' \leq T} S_{T'}$ is of order $\sqrt{T \ln \ln T}$.²⁰ Because $\frac{\sqrt{T \ln \ln T}}{\sqrt{T \ln T}}$ goes to 0 as T grows large, hurdle Θ_T insures that uninformed agents have very little hope to obtain unjustified returns. Indeed, the following result holds.

Lemma 3 (hurdle effectiveness). *If the agent is uninformed, then for any allocation strategy a , any environment \mathcal{P} and any horizon N ,*

$$\mathbb{E}_a \left(\sum_{t=1}^N \mathbf{1}_{S_t \geq \Theta_t} \right) \leq \frac{\pi^2}{2} \exp(-2M),$$

where π is the constant 3.1415...

Because hurdles also reduce the payoffs accruing to informed agents, they carry an incentive cost. Still as the next theorem shows, this incentive cost is moderate. Denote by r_{λ, π^Θ} the net expected per-period returns generated by the agent under the hurdled calibrated

²⁰See Billingsley (1995), Theorem 9.5.

contract. The following result holds.

Theorem 4 (performance with screening). *Pick $\alpha_0 \in (0, 1)$ and for any $\eta \in (0, 1)$, let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists a constant m independent of time horizon N and probability space \mathcal{P} such that for all h_T ,*

$$r_{\lambda, \pi^\Theta} |h_T \geq (1 - \eta)r_{\alpha_0} |h_T - m \sqrt{\frac{\ln N}{N}} \quad (26)$$

Furthermore, whenever $-b + \alpha w \bar{d} \times \frac{\pi^2}{2} \exp(-2M) < 0$, it is strictly optimal for uninformed agents not to participate.

The combination of initial fee $-b$ and hurdle Θ_t induces self-screening by uninformed agents. Hurdle Θ_t is large enough that uninformed agents have little hope to be rewarded by luck but small enough that it does not significantly affect the incentives of informed agents. The penalty which was of order $\frac{1}{\sqrt{N}}$ in Theorem 2 is now of order $\sqrt{\frac{\ln N}{N}}$. It is possible to prove an improved bound in some circumstances. Appendix B shows that when returns are grainy, i.e. either zero or bounded away from 0, the performance loss from screening is only of order $\frac{1}{\sqrt{N}}$.

5.3 Multi-agent contracts

The analysis so far has focused on contracting with a single agent. This section shows how to extend the logic of Sections 3 and 4 to environments with multiple agents. The framework is identical to that of Section 2 except that there are now K agents denoted by $k \in \{1, \dots, K\}$, each of whom makes private information acquisition decisions $c_{k,t} \in [0, +\infty)$, inducing a filtration \mathcal{F}_t^k . In each period t , manager k suggests an asset allocation $a_{k,t}$ inducing potential wealth $w_{k,t} = w(1 + \langle a_{k,t}, \mathbf{r}_t \rangle)$. As in Section 2 the environment is general. Public and private signals $(I_t^0, I_c^k(c_{k,t}))_{k \in \{1, \dots, K\}}$ are arbitrary random variables from an underlying measurable state space (Ω, σ) to a measurable signal space $(\mathcal{I}, \sigma_{\mathcal{I}})$. The environment $\mathcal{P} = (\Omega, \sigma, P)$ is specified by defining a probability measure P on (Ω, σ) . This probability measure is

unrestricted: the agents may have access to different information, their respective ability to generate information may differ and vary over time, or be correlated in arbitrary ways. Filtration \mathcal{F}_t^0 still denotes the public information filtration available to the principal. The first step of the analysis extends the high-liability benchmark contract of Section 3. The second step of the analysis shows how to calibrate this high-liability contract.

Multi-agent benchmark contracts. The multi-agent contract described here is a direct extension of the linear contract described in Section 3. Each agent $k \geq 1$ is paid according to a linear contract in which the allocation of agent $k - 1$ serves as the default allocation previously corresponding to a_t^0 .

Specifically, in each period t , allocations $a_{k,t}$ are submitted by agents increasing order of k . This ordering is a constraint imposed by the mechanism. The mechanism informs each agent k of the allocations $(a_{k',t})_{k' < k}$ chosen by agents $k' < k$. Agent k receives no information about the allocations chosen by agents $k'' > k$. Under the benchmark contract, payments $\pi_{k,t}$ to agent k are defined by

$$\forall k \in \{1, \dots, K\}, \quad \pi_{k,t} = \alpha(w_{k,t} - w_{k-1,t}). \quad (27)$$

Finally, resources are invested according to the allocation $a_{K,t}$ suggested by the last agent. Under this multi-agent contract, each agent is paid a share α of his externality on the principal, taking into account the information provided by previous agents.

The strategy profile (c_k, a_k) of agent k must be adapted to the information available to the agent (by construction this includes allocations by previous managers). The set of such adapted strategies is denoted by $\mathcal{C}_k \times \mathcal{A}_k$.²¹ Furthermore define $(c, a) = (c_k, a_k)_{k \in \{1, \dots, K\}}$ and $\mathcal{C} \times \mathcal{A} = \prod_{k \in \{1, \dots, K\}} \mathcal{C}_k \times \mathcal{A}_k$ the set of adapted strategy profiles. For any $\hat{c} \in [0, +\infty)$, the

²¹ Because of the hierarchical structure of the mechanism, agent $k' < k$ is indifferent about whether or not to send information to agent k . Agent k' could be made to strictly prefer sharing information by being awarded a small share of manager k 's profits. The analysis that follows holds for any amount of information provided by previous managers to future managers. Different assumptions about such information transmission simply correspond to different measurability constraints on the class of strategies $\mathcal{C}_k \times \mathcal{A}_k$.

maximum returns that can be obtained at an expected per-period cost of \hat{c} are denoted by

$$r_{\max}(\hat{c}) = \max_{\substack{(c,a) \in \mathcal{C} \times \mathcal{A} \\ \frac{1}{N} \mathbb{E}(\sum_{k,t} c_{k,t}) \leq \hat{c}}} \frac{1}{wN} \mathbb{E}_{c,a} \left(\sum_{t=1}^N w_t^K - w_t^0 \right).^{22}$$

Denote by r_α the average returns accruing to the principal under this benchmark contract. The following bound extends point (i) of Theorem 1.

Lemma 4. *For any environment \mathcal{P} ,*

$$r_\alpha \geq (1 - \alpha) \max_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right).$$

As in Theorem 1, this lower bound implies that by optimizing over α , one can find a benchmark contract that guarantees a fixed share of the first-best surplus. Similarly to the benchmark contract of Section 3, this contract is weakly renegotiation proof and satisfies no-loss.

Calibrated contracts. The high-liability multi-agent contract described in (27) can be calibrated using the methods of Section 4. The main difference is that there are now K scaling factors $\lambda_t = (\lambda_{k,t})_{k \in \{1, \dots, K\}} \in [0, 1]^K$ that are used to define adjusted allocations $a_{k,t}^\lambda$ in the following recursive manner:

$$a_{0,t}^\lambda = a_{0,t} \quad \text{and} \quad \forall k \geq 1, \quad a_{k,t}^\lambda = \lambda_{k,t} a_{k,t} + (1 - \lambda_{k,t}) a_{k-1,t}^\lambda.$$

²²Note that these returns depend on the measurability constraints imposed on $\mathcal{C} \times \mathcal{A}$ (see footnote 21).

Let $w_{k,t}^\lambda$ denote the corresponding wealth realizations. For all $k \geq 1$, define regrets

$$\mathcal{R}_{k,T}^1 = \sum_{t=0}^T \pi_t - \alpha(w_{k,t}^\lambda - w_{k-1,t}^\lambda) \quad (\text{no-excess payments}) \quad (28)$$

$$\mathcal{R}_{k,T}^2 = \max_{T' \leq T} \sum_{t=T'}^T w_{k,t} - w_{k,t}^\lambda \quad (\text{no foregone returns}). \quad (29)$$

Keeping these regrets small corresponds to implementing appropriate generalizations of incentive properties (14) and (15) for all agents. The usual approachability condition implies that regrets $(\mathcal{R}_{k,T}^1, \mathcal{R}_{k,T}^2)_{k \in \{1, \dots, K\}}$ can be kept small by choosing contract parameters $(\lambda_k, \pi_k)_{k \in K}$ according to,

$$\lambda_{k,T+1} = \frac{\alpha [\mathcal{R}_{k,T}^2]^+}{\alpha [\mathcal{R}_{k,T}^2]^+ + [\mathcal{R}_{k,T}^1]^+} \quad \text{and} \quad \pi_{k,T+1} = \begin{cases} \alpha(w_{k,T+1}^\lambda - w_{k-1,T+1}^\lambda)^+ & \text{if } \mathcal{R}_{k,T}^1 \leq 0 \\ 0 & \text{otherwise} \end{cases}.$$

Under this calibrated multi-agent contract the following result obtains.

Theorem 5. *Pick $\alpha_0 > 0$ and for $\eta \in (0, 1)$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists a constant m independent of environment \mathcal{P} , time horizon N and number of agents K such that the multi-agent calibrated contract $(\lambda, \pi) = (\pi_k, \lambda_k)_{k \in K}$ of parameter α satisfies*

$$r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - m\sqrt{\frac{K}{N}} \quad (30)$$

$$\forall h_T, \quad r_{\lambda, \pi}|h_T \geq (1 - \eta)r_{\alpha_0}|h_T - m\sqrt{\frac{K}{N}}. \quad (31)$$

This extends the approach of Sections 3 and 4 to environments where the principal relies on the information acquired by several agents. Note that the bounds provided by Theorem 5 are useful only if K is small compared to N . In this respect, inducing uninformed agent to self-screen will help reduce K . The screening strategy developed in Section 5.2 can be adapted without difficulty to multi-agent contracts.

6 Discussion

This paper develops a robust approach to dynamic contracting in two steps. The first step identifies high-liability linear contracts that satisfy attractive efficiency properties regardless of the underlying environment. The second step shows how to approximate the performance of benchmark contracts using limited-liability dynamic contracts. The contracting strategy is to calibrate rewards to the agent as well as the share of wealth he manages so that key properties of the benchmark contract are approximately replicated. The resulting calibrated contracts are simple, and perform approximately as well as an attractive benchmark under very general conditions. The approach can be extended to calibrate a much broader class of high-liability contracts, induce self-screening by uninformed agents, or jointly manage several agents. Appendix A shows that the analysis is robust to partial departures from rationality by the agent and can be extended to the case where future payoffs are discounted.

The remainder of this section discusses in further detail how calibrated contracts relate to other contracts of interest and delineates possible avenues for future research.

6.1 Relation to Other Contracts

High-watermark contracts. The calibrated contracts described in Section 4 are closely related to the high-watermark contracts frequently used in the financial industry (see for instance Goetzmann et al. (2003) who develop an option-pricing approach to high-watermark contracts, or Panageas and Westerfield (2009) who show in a specific context that high-watermark contracts need not lead to excessive risk-taking). High-watermark contracts are structured as follows: at time T , the investment share λ_T is always 1, and the agent gets paid

$$\pi_T^{\text{wmk}} = \alpha \left(\sum_{t=1}^T w_t - w_t^0 - \max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right] \right)^+. \quad (32)$$

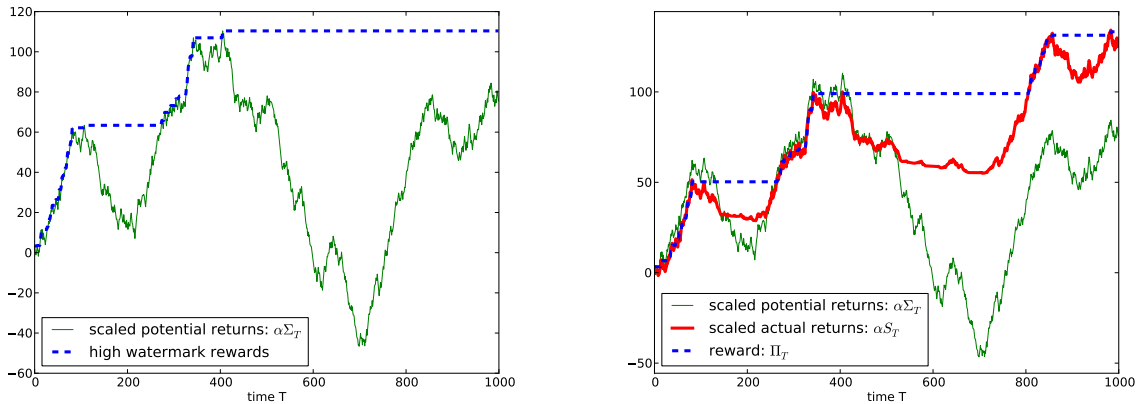
Quantity $\max_{T' < T} \left[\sum_{t=1}^{T'} w_t - w_t^0 \right]$ is referred to as the high-watermark and represents the maximum historical cumulated returns at time T . The agent only gets paid when he improves on his own historical performance. Note that high-watermark contracts are dynamic and satisfy limited liability constraint (4). In particular, for all T , $\pi_T \in [0, w_T]$.

High-watermark contracts, as well as calibrated contracts, attempt to reward the agent a share α of his externality on the principal. In other words, both types of contracts attempt to keep aggregate rewards Π_T close to αS_T . Lemma 1 shows that calibrated contracts achieve this goal for any realization of uncertainty and any allocation strategy. High-watermark keep aggregate rewards Π_T close to αS_T along paths such that the process for value-added $(S_T)_{T \geq 1}$ is on average increasing, but not if the process for value-added has significant downturns. This is illustrated by Figure 1(a). Whenever there is an extended drop in performance, the relationship between rewards Π_T and performance αS_T breaks down. Indeed Π_T is by construction weakly increasing while under the high-watermark contract, αS_T can decrease in arbitrary ways. More explicitly, imagine that the agent delivers performance $(1, 1, -1, 1, 1, -1, \dots, 1, 1, -1)$ so that total surplus is $N/3$. Cumulated value-added S_T is on average increasing and under the high-watermark contract the agent obtains a reward $\alpha N/3 + o(N)$. If instead the agent delivers returns $(1, 1, \dots, 1)$ for the first $2N/3$ periods, followed by $(-1, -1, \dots, -1)$ for the last $N/3$ periods, then the surplus generated by the agent is still $N/3$, but the high-watermark contract now gives him a payoff $2\alpha N/3 + o(N)$.

This has two implications. First, because the agent does not suffer from extended downturns, an agent who has lost the ability to generate positive return (e.g. his information has become unreliable) may cause large losses by choosing negative expected value investments that generate variance. Second, if a talented agent has been unlucky and experienced a drop in returns, the difficulty of catching up with a high watermark may discourage investment altogether. As a result high-watermark contracts exhibit large gains to renegotiation. If a manager performed well for an extended amount of time, following which he experiences sharp losses, the principal and the agent may both benefit strongly from forgiving the losses

and pretending that the current high-watermark is lower than it really is.

As Figure 1(b) illustrates, calibrated contracts ultimately boil down to writing a high-watermark contract on the modified performance measure S_T . By choosing investment shares λ_t appropriately, calibrated contracts are able to keep tight the relationship between rewards Π_T and actual performance S_T at every history. As a result, extended downturns have a much more limited impact on incentives. Note that investment shares λ_t must move smoothly with performance instead of taking values 0 or 1. Rather than a stop-loss provision, it is more accurate to think of the calibrated investment shares $(\lambda_t)_{t \geq 0}$ as continuously implementing a robust option on the agent’s potential performance Σ_T .²³ The fact that calibrated contracts do not generate large foregone performance (Lemma 1) implies that along parts of the path of play where the agent is generating positive returns, investment shares will be close to one. Inversely, investment shares may be significantly below one along portions of the path where the agent is not generating positive returns. As discussed in Appendix A this makes calibrated contracts robust to temporarily suboptimal play by the agent.



(a) high-watermark contract: potential returns $\alpha\Sigma_T$, (b) calibrated contracts: potential returns $\alpha\Sigma_T$, actual returns αS_T , rewards Π_T

Figure 1: high-watermark and calibrated contracts for a sample path of potential returns $(\Sigma_T)_{T \geq 1}$, with target reward rate $\alpha = 20\%$.

²³See DeMarzo et al. (2009) for work on the relation between approachability and robust option pricing.

Connection with optimal contracting. It is instructive to note that for sufficiently high discount factors, DeMarzo and Sannikov (2006) as well as Biais et al. (2007, 2010) derive high-watermark contracts as optimal contracts in their specific environments. The link is not entirely obvious because their optimal contracts are described in the standard (forward looking) language of continuation values. Because calibrated (and high-watermark) contracts are detail-free, they can only be described in reference to (backward looking) realized observables. This difference however is superficial and the connection between the two approaches is in fact significant.²⁴ To a first order, DeMarzo and Sannikov (2006) and Biais et al. (2007, 2010) find that in their environment, under optimal contracts, the agent’s continuation value follows a Brownian motion, proportional to the agent’s performance, and reflected at some upper bound \bar{W} . Whenever the agent’s continuation value hits this upper bound, he is paid a fixed proportion of the surplus he generates. This in fact encodes for a high-watermark contract. Imagine that at time t , the agent is promised value \bar{W} , and that he starts losing money. Then, his continuation value moves in a way proportional to his performance, and he is only paid again when his performance covers his losses so that his continuation value climbs back to \bar{W} . This coincides with the reward profile of a high-watermark contract: the agent only gets paid once he has recouped his losses. This connection should not be entirely surprising: DeMarzo and Sannikov (2006) as well as Biais et al. (2007, 2010) consider environments with linear production technology in which the benchmark high-liability contracts of Section 3 are close to optimal; calibrated contracts are specifically designed to approximate the performance of such contracts.

The connection is particularly strong with Biais et al. (2007) and especially Biais et al. (2010) who emphasize the role of downsizing the size of the project managed by the agent as a function of his performance. This is related to varying investment shares $(\lambda_t)_{t \geq 1}$ in the current paper. The use of downsizing in Biais et al. (2007, 2010) however is slightly different.

²⁴The optimal contracts derived by DeMarzo and Sannikov (2006) and Biais et al. (2007, 2010) can be given a backward looking description since there is a one to one mapping between realized payoffs and continuation values.

In their work, downsizing occurs when continuation values are so low that at the current size of the project, optimal behavior can no longer be enforced. Downsizing allows to deliver the promised low values while maintaining appropriate incentive compatibility conditions in the continuation game. As a result, downsizing occurs only after sufficiently long strings of poor performance. In the current paper, $(\lambda_t)_{t \geq 1}$ can be seen as a *preventive* downsizing scheme, which rules out continuation values so low that incentive provision becomes problematic.

6.2 Future Work

The relative simplicity of the analysis presented in the paper gives reasonable hope that the approach may be gainfully used in other settings. Three directions seem particularly attractive for further theoretical work. First, it would be valuable to develop a better understanding of how more restrictive max min problems map into different benchmark contracts, and whether these contracts can be calibrated. A second challenge is to allow for risk-aversion. Some suggestions are offered in Chassang (2011), but more work remains to be done, part of the difficulty being to characterize appropriate benchmark contracts. A third avenue for research is to expand the analysis of multi-player environments. This may be useful to fine tune parameters of the contract such as reward rate α , or select the most promising K agents out of a large number. This may also be interesting beyond the principal-agent setting that this paper focuses on. For instance many attractive allocation mechanisms, such as Vickrey-Clarke-Groves mechanisms, require agents to make significant payments and are therefore ill-suited in environments where agents are severely cash constrained. A dynamic calibration approach such as the one developed in this paper may help relax such limited liability constraints.

In addition, with actual implementation in mind, it seems important to determine whether calibrated contracts really do induce approximately good behavior from agents. Indeed, Theorem 2 relies on the agent's ability to understand the dynamic incentive properties of calibrated contracts. This is a demanding rationality requirement and whether or not it

holds in practice is ultimately an empirical question. An advantage of the detail free approach developed here is that it lends itself naturally to realistic experiments using actual returns data, since the contracts should perform well regardless of the agent’s beliefs over the process for returns.

Appendix

A Additional Results

This appendix presents a number of additional results. Appendix A.1 shows that the calibrated contracts of Section 4 perform well even if the agent isn’t rational and behaves suboptimally over any arbitrary interval of time. Appendix A.2 extends the analysis to the case where principal and agent discount future payoffs; Appendix A.3 allows for varying levels of wealth. Appendix A.4 considers principal-agent problems where the set of actions is not convex.

A.1 Robustness to Accidents

The analysis presented in this paper assumes that the agent is rational. It turns out that calibrated contracts are robust to the possibility of “accidents” during which the agent can behave sub-optimally over an extended amount of time.

An accident may correspond to a temporary mistake in the agent’s trading strategy or an error in his data; alternatively, the agent may be temporarily irrational or have unmodeled incentives to misbehave (e.g. he is bribed to unload bad risks on the principal). Formally, this is modeled by assuming that during a random time interval $[T_1, T_2]$ —in the accident state—the agent uses an exogenously specified allocation strategy a_t^Δ .²⁵ This strategy may be arbitrarily bad (within the bounds imposed by Assumption 1) and need only be measurable

²⁵The analysis given here allows accidents to occur over a single interval of time. The analysis extends without change to environments with a bounded number of intervals.

with respect to \mathcal{F}_N . For instance, during the lapse of the accident, the agent could pick the worst ex post asset allocation in every period. Robustness to accidents of this kind is related to Eliaz (2002) which studies how well mechanisms perform if some players are faulty, i.e. if they use non-optimal strategies. Here, robustness to accidents corresponds to fault tolerance with respect to the agent's selves over $[T_1, T_2]$.

It should be noted that in this environment, the benchmark linear contract is no longer sufficient to guarantee good performance. Accidents can undo all the profit generated by the well incentivized agent in his normal state. Strikingly, in spite of accidents, calibrated contracts are such that the excess returns generated by the agent will be approximately as high as the returns he could generate when accidents are “lucky”, i.e. when the exogenous allocation during accident states is

$$\forall T \in [T_1, T_2], \quad a_T^{\Delta\Delta} = \begin{cases} a_T^0 & \text{if } \sum_{t=T_1}^{T_2} w_t^\Delta - w_t^0 < 0 \quad (\text{accident is unlucky}) \\ a_T^\Delta & \text{if } \sum_{t=T_1}^{T_2} w_t^\Delta - w_t^0 > 0 \quad (\text{accident is lucky}) \end{cases}$$

where w_t^Δ is the realized wealth under the a_t^Δ at time t . Denote by $r_{\lambda, \pi}^\Delta$ the net expected returns to the principal when accidental behavior is a_t^Δ and the calibrated contract is used. Denote by $r_\alpha^{\Delta\Delta}$ the net expected returns to the principal when accidental behavior is $a_t^{\Delta\Delta}$ and the benchmark contract of parameter α is used. Consider the contract (λ, π) defined by (10) and (11). The following holds.

Theorem A.1 (accident proofness). *Pick α_0 and for any $\eta > 0$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists a constant m , independent of N and \mathcal{P} such that,*

$$r_{\lambda, \pi}^\Delta \geq (1 - \eta)r_{\alpha_0}^{\Delta\Delta} - \frac{m}{\sqrt{N}}.$$

Proof. Let $w_t^{\Delta\Delta}$ and $\Sigma_N^{\Delta\Delta} = \sum_{t=1}^N w_t^{\Delta\Delta} - w_t^0$ denote potential realized wealth and aggregate excess returns when accidents are lucky. The notation of Section 4 extends, adding superscripts $^\Delta$ and $^{\Delta\Delta}$ to denote relevant objects under the original accidental allocation a^Δ , and

under the lucky accidental allocation $a^{\Delta\Delta}$. The key step is to provide an adequate extension of Lemma 1.

Inequality (18) still applies, and we necessarily have that

$$-\alpha w\bar{d} \leq \Pi_N^\Delta - \alpha S_N^\Delta \leq \alpha w\bar{d}\sqrt{N}. \quad (33)$$

In addition, let us show that for any investment strategy of the agent,

$$\Sigma_N^{\Delta\Delta} - 4w\bar{d}\sqrt{N} \leq S_N^\Delta \quad (34)$$

i.e. up to an order \sqrt{N} , given any investment strategy, the actual excess returns generated under the calibrated contract are at least as high as the returns generated when accidents are lucky. We have that $\Sigma_N^{\Delta\Delta} = \Sigma_{N \setminus T_2+1}^\Delta + \left[\Sigma_{T_2 \setminus T_1}^\Delta\right]^+ + \Sigma_{T_1-1}^\Delta$. Because inequality (17) still holds, this implies that

$$\Sigma_N^{\Delta\Delta} \leq \begin{cases} S_N^\Delta + w\bar{d}\sqrt{N} & \text{if } \Sigma_{T_2 \setminus T_1}^\Delta > 0 \\ S_{N \setminus T_2+1}^\Delta + S_{T_1-1}^\Delta + 3w\bar{d}\sqrt{N} & \text{otherwise} \end{cases}$$

By (33), it follows that

$$\begin{aligned} \Pi_{T_2}^\Delta - \alpha w\bar{d}\sqrt{T_2} &\leq \alpha S_{T_2}^\Delta \leq \Pi_{T_2}^\Delta + \alpha w\bar{d} \\ \Pi_{T_1-1}^\Delta - \alpha w\bar{d}\sqrt{T_1-1} &\leq \alpha S_{T_1-1}^\Delta \leq \Pi_{T_1-1}^\Delta + \alpha w\bar{d}. \end{aligned}$$

Subtracting these two inequalities yields that,

$$\Pi_{T_2 \setminus T_1}^\Delta - \alpha w\bar{d}(1 + \sqrt{T_2}) \leq \alpha(S_{T_2}^\Delta - S_{T_1-1}^\Delta) = \alpha(S_{T_2 \setminus T_1}^\Delta).$$

Since flow rewards are weakly positive, $\Pi_{T_2 \setminus T_1}^\Delta \geq 0$, which implies that for any realization of

returns,

$$\begin{aligned}\Sigma_N^{\Delta\Delta} &\leq S_{N \setminus T_2+1}^\Delta + S_{T_2 \setminus T_1}^\Delta + S_{T_1-1}^\Delta + 4w\bar{d}\sqrt{N} \\ &\leq S_N^\Delta + 4w\bar{d}\sqrt{N}.\end{aligned}$$

This implies (34). Given (33) and (34). Theorem A.1 follows from a reasoning identical to that of Lemma 2. \square

A.2 Discounting

The analysis of Section 4 can be extended to environments where principal and agent discount the future by a factor δ so that the agent's payoffs are $\mathbb{E}\left(\sum_{t=1}^N \delta^{t-1}(\pi_t - c_t)\right)$ and the principal's surplus is $\mathbb{E}\left(\sum_{t=1}^N \delta^{t-1}(w_t - w_t^0)\right)$. Let $N_\delta = \sum_{t=1}^N \delta^t$. This appendix shows that under discounting, the performance bound of Theorem 2 extends with a loss of order $\sqrt{1/N_\delta}$ instead of $\sqrt{1/N}$.

Benchmark contract. The benchmark contract is still to reward the agent $\alpha(w_t - w_t^0)$ in every period t . This linear contract guarantees the principal a robust payoff bound similar to that given for the benchmark contract in Section 3. For any contract (λ, π) , where λ may be constant and equal to 1, define

$$r_{\lambda, \pi} = \inf \left\{ \mathbb{E}_{c, a} \left(\frac{1}{wN_\delta} \sum_{t=1}^N \delta^{t-1} [\lambda_t(w_t - w_t^0) - \pi_t] \right) \middle| (c, a) \text{ solves } \max_{c, a} \mathbb{E} \left(\sum_{t=1}^N \delta^{t-1} [\pi_t - c_t] \right) \right\}$$

the average discounted per-period returns accruing to the principal under contract (λ, π) . Let r_α denote returns accruing to the principal under the benchmark contract. In addition define

$$r_{\max}(\hat{c}) \equiv \sup_{\substack{c \text{ s.t.} \\ \mathbb{E} \left[\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} c_t \right] \leq \hat{c}}} \mathbb{E}_{c, a^*} \left(\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} \langle a_t^* - a_t^0, r_t \rangle \right)$$

the maximum discounted per-period returns that can be generated at an expected discounted per-period cost of \hat{c} .

Lemma A.1. *For all environments \mathcal{P} ,*

$$r_\alpha \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} \right).$$

Proof. The proof is identical to that of Theorem 1, point (i). □

Calibration. The calibrated contract is built using the following regrets

$$\mathcal{R}_{1,T} = \sum_{t=1}^T \delta^{t-1} (\pi_t - \alpha(w_t - w_t^0)) \quad \text{and} \quad \mathcal{R}_{2,T} = \max_{T \leq T'} \sum_{t=T'}^T \delta^{t-1} (1 - \lambda_t)(w_t - w_t^0).$$

Contract parameters $(\lambda_t, \pi_t)_{t \geq 1}$ are computed recursively according to

$$\lambda_t = \frac{\alpha \mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha \mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_t = \begin{cases} \alpha(w_t - w_t^0) & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise} \end{cases}.$$

The following result extends Lemma 1, showing that incentives are approximately correct.

Lemma A.2 (approximate incentives). *For all T , and all possible histories,*

$$\frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} (1 - \lambda_t)(w_t - w_t^0) \leq \frac{w\bar{d}}{\sqrt{N_\delta}} \tag{35}$$

$$-\frac{w\bar{d}}{N_\delta} \leq \frac{1}{N_\delta} \sum_{t=1}^N \delta^{t-1} [\pi_t - \alpha(w_t - w_t^0)] \leq \frac{w\bar{d}}{\sqrt{N_\delta}}. \tag{36}$$

Proof. Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \alpha \mathcal{R}_{2,T})$ denote the vector of regrets, and $\rho_{T+1} = \mathcal{R}_{T+1} - \mathcal{R}_T$. Contract (λ, π) is calibrated so that in every period $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$. It follows that

$$\|\mathcal{R}_N^+\|^2 \leq \sum_{t=1}^N \|\rho_T\|^2.$$

Furthermore, we have that $\|\rho_T\|^2 \leq \delta^{2T} w \bar{d}$, which implies that

$$\|\mathcal{R}_T^+\|^2 \leq w \bar{d} \sum_{t=1}^N \delta^{2(t-1)} \leq w \bar{d} \sum_{t=1}^N \delta^{t-1}.$$

This implies the right-hand sides of (35) and (36). The left-hand side of (36) follows from a proof identical to that of the left-hand side of (18). \square

This implies the following bounds for returns $r_{\lambda, \pi}$.

Theorem A.2. *Pick α_0 and for $\eta > 0$, let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists $m \geq 0$ such that for all environments \mathcal{P} , all δ and all N ,*

$$r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - \frac{m}{\sqrt{N_\delta}} \quad (37)$$

$$r_{\lambda, \pi} \geq (1 - \alpha) \sup_{\hat{c} \in [0, +\infty)} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} - \frac{3\bar{d}}{\sqrt{N_\delta}} \right). \quad (38)$$

Inequality (37) is an extension of (12) but involves an unspecified constant m which may be large in practice. Inequality (38) gives a more explicit performance bound directly related to that given in Lemma A.1 for the benchmark contract. Note that in both cases the proof does not attempt to optimize constants and improved bounds can be obtained.²⁶

Proof. Inequality (37) follows from Lemma A.2 and Lemma 2.

Inequality (38) follows from Lemma A.2 and a reasoning identical to the proof of Theorem 1, point (i). Denote by (c, a) the optimal strategy profile under contract (λ, π) . Under any alternative strategy profile (c', a') , we have

$$\mathbb{E}_{c, a} \left(\sum_{t=1}^N \delta^{t-1} (\pi_t - c_t) \right) \geq \mathbb{E}_{c', a'} \left(\sum_{t=1}^N \delta^{t-1} (\pi_t - c'_t) \right)$$

²⁶For instance the constant 3 in (38) can be replaced by $2 + 1/\sqrt{N_\delta}$.

Using (36) and (35), this implies that

$$\mathbb{E}_{c,a} \left(\sum_{t=1}^N \delta^{t-1} [\alpha \lambda_t (w_t - w_t^0) - c_t] \right) + w\bar{d}\sqrt{N_\delta} \geq \mathbb{E}_{c',a'} \left(\sum_{t=1}^N \delta^{t-1} [\alpha (w_t - w_t^0) - c_t] \right) - w\bar{d} \left(\sqrt{N_\delta} + 1 \right)$$

Since this holds for every strategy profile (c', a') , it follows from rearranging that,

$$r_{\lambda,\pi} \geq (1 - \alpha) \sup_{\hat{c} \geq 0} \left(r_{\max}(\hat{c}) - \frac{\hat{c}}{\alpha w} - \bar{d} \frac{2\sqrt{N_\delta} + 1}{N_\delta} \right).$$

□

A.3 Varying Wealth

The calibrated contracts described in Section 4 perform equally well if the invested wealth in each period varies within some set $[0, \bar{w}]$. Let w_t^i denote the initial invested wealth in period t . Given a contract (λ, π) , quantities Σ_T, S_T and Π_T are defined as

$$\Pi_T = \sum_{t=1}^T \pi_t; \quad \Sigma_T = \sum_{t=1}^T w_t^i \langle a_t - a_t^0, \mathbf{r}_t \rangle; \quad S_T = \sum_{t=1}^T \lambda_t w_t^i \langle a_t - a_t^0, \mathbf{r}_t \rangle.$$

Similarly, let $\Sigma_{T \setminus T'} = \Sigma_T - \Sigma_{T'-1}$, $\Pi_{T \setminus T'} = \Pi_T - \Pi_{T'-1}$, $S_{T \setminus T'} = S_T - S_{T'-1}$. As in Section 4, regrets $\mathcal{R}_{1,T}$ and $\mathcal{R}_{2,T}$ are defined by

$$\mathcal{R}_{1,T} = \Pi_T - \alpha S_T \quad \text{and} \quad \mathcal{R}_{2,T} = \max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}.$$

Contract (λ, π) is unchanged:

$$\lambda_{T+1} = \frac{\alpha \mathcal{R}_{2,T}^+}{\alpha \mathcal{R}_{2,T}^+ + \mathcal{R}_{1,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} \alpha \lambda_{T+1} (w_{T+1} - w_{T+1}^0)^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \\ 0 & \text{otherwise} \end{cases}.$$

Under this adjusted contract, Theorem 2 extends as is, with an identical proof.

A.4 Principal-Agent Problems Without a Convex Action Space

This appendix extends the analysis to a principal-agent framework more general than the financial contracting problem studied in the paper. The principal and the agent are still risk-neutral, but in every period the agent suggests and implements an action $a_t \in A$, where A is a potentially non-convex set of actions. Every period, a state of the world \mathbf{r}_t is drawn which, given action a , yields observable payoffs $w(a, \mathbf{r}_t)$ to the principal. Cost c_t may now represent the cost of information acquisition or the cost of taking a specific action. Action a_t^0 is the action that the principal would (could) implement on her own. The main difference is that because set A need not be convex, the principal must use randomized strategies to calibrate her contract with the agent.

The calibrated contract of Section 4 can be adapted as follows. Parameter λ_t now denotes the probability that the principal follows or authorizes the action suggested by the agent.²⁷ Let a_t^λ denote the action actually taken at time t . Denote by $\psi_t \equiv w(a_t, \mathbf{r}_t) - w(a_t^0, \mathbf{r}_t)$ the potential excess returns and by $\psi_t^\lambda \equiv w(a_t^\lambda, \mathbf{r}_t) - w(a_t^0, \mathbf{r}_t)$ the realized excess returns. The benchmark contract is to reward the agent $\alpha\psi_t$ in each period. As in Section 4, define

$$\Sigma_T = \sum_{t=1}^T \psi_t, \quad \Pi_T = \sum_{t=1}^T \pi_t, \quad S_T = \sum_{t=1}^T \psi_t^\lambda,$$

as well as $\Sigma_{T \setminus T'} = \Sigma_T - \Sigma_{T'-1}$, $\Pi_{T \setminus T'} = \Pi_T - \Pi_{T'-1}$ and $S_{T \setminus T'} = S_T - S_{T'-1}$. As in Section 4, regrets are defined by

$$\mathcal{R}_{1,T} \equiv \Pi_T - \alpha S_T \quad \text{and} \quad \mathcal{R}_{2,T} \equiv \max_{T' \leq T} \Sigma_{T \setminus T'} - S_{T \setminus T'}.$$

²⁷For calibration results to hold, it is important that the agent not be able to condition his suggested action on the outcome of the principal's randomization. If the agent takes the action on behalf of the principal, λ_t should be interpreted as the probability that the principal approve the agent's proposed course of action.

Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \alpha\mathcal{R}_{2,T})$. Contract $(\lambda_t, \pi_t)_{t \in \mathbb{N}}$ is defined by

$$\lambda_{T+1} = \frac{\alpha\mathcal{R}_{2,T}^+}{\mathcal{R}_{1,T}^+ + \alpha\mathcal{R}_{2,T}^+} \quad \text{and} \quad \pi_{T+1} = \begin{cases} 0 & \text{if } \mathcal{R}_{1,T} > 0 \\ \alpha [\psi_{T+1}^\lambda]^+ & \text{if } \mathcal{R}_{1,T} \leq 0 \end{cases} \quad (39)$$

with the convention that $\frac{0}{0} = 1$. Lemma 1 extends as follows.

Lemma A.3 (approximate incentives). *For all T , and any strategy (c, a) of the agent, we have that*

$$\mathbb{E}_{c,a}\Sigma_T - \mathbb{E}_{c,a}S_T \leq w\bar{d}\sqrt{T} \quad (40)$$

$$-\alpha w\bar{d} \leq \Pi_T - \alpha\mathbb{E}_{c,a}S_T \leq \alpha w\bar{d}\sqrt{T}. \quad (41)$$

Proof. The left-hand side of (41) follows from a proof identical to that of the left-hand side of (18).

Let us turn to the other inequalities. Let $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$ denote flow regrets, and observe that $\mathbb{E}_{c,a}(\langle \mathcal{R}_{T-1}^+, \rho_T \rangle) \leq 0$. Hence, a proof identical to that of Lemma 1 yields that

$$\forall(c, a), \quad \mathbb{E}_{c,a} \|\mathcal{R}_T^+\|^2 \leq (\alpha\bar{d}w)^2 T.$$

It follows from Jensen's inequality that for all $i \in \{1, 2\}$,

$$\mathbb{E}_{c,a}(\mathcal{R}_{i,t}^+) \leq \mathbb{E}_{c,a} \left(\sqrt{[\mathcal{R}_{i,t}^+]^2} \right) \leq \sqrt{\mathbb{E}_{c,a}([\mathcal{R}_{i,t}^+]^2)} \leq \alpha w\bar{d}\sqrt{T}.$$

This implies (40) and the right-hand side of (41). \square

Given Lemma A.3, a proof identical to that of Lemma 2 yields the following performance bound: pick $\alpha_0, \eta > 0$ and let $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. There exists m independent of N and \mathcal{P} such that

$$r_{\lambda,\pi} \geq (1 - \eta)r_{\alpha_0} - \frac{m}{\sqrt{N}}.$$

B Proofs

B.1 Proofs for Section 3

Proof of Theorem 1: Let us begin with point (i). Let (c_α, a^*) denote the agent's policy under the benchmark contract. For any $\hat{c} \in [0, +\infty)$, denote by (c, a^*) a surplus maximizing policy with average expected cost $\frac{1}{N}\mathbb{E}\left(\sum_{t=1}^N c_t\right) \leq \hat{c}$. Recall that a^* denotes the optimal allocation strategy conditional on information. Since policy (c, a^*) guarantees the agent a per-period payoff of $\alpha wr_{\max}(\hat{c}) - \hat{c}$, it must be that $\frac{\alpha}{1-\alpha}wr_\alpha - \mathbb{E}c_\alpha \geq \alpha wr_{\max}(\hat{c}) - \hat{c}$. Since the agent must expend weakly positive effort, this implies that $\frac{\alpha}{1-\alpha}wr_\alpha \geq \alpha wr_{\max}(\hat{c}) - \hat{c}$, which yields point (i).

Let us now turn to point (ii). Using the bound given in point (i) with $\alpha = \sqrt{\rho}$ and $\hat{c} = c_{FB}$ yields after manipulation that

$$\frac{wr_\alpha}{wr_{FB} - c_{FB}} \geq 1 - 2\frac{\sqrt{\rho}}{1 + \sqrt{\rho}}.$$

We now show that no contract can improve on this bound over the class of environments \mathbb{P}_ρ . For this it is sufficient to show that no contract can improve on this bound for some subclass of environments included in \mathbb{P}_ρ . We consider the following family of settings.

There are two assets, 1 and 2. Asset 1 is riskless with returns $r_{1,t} = 0$ every period. Asset 2 is risky and i.i.d. with negative expected value. Specifically, $r_{2,t} = 1$ with probability $1/3$ and $r_{2,t} = -1$ with probability $2/3$. The agent can only acquire information in period $t = 1$, but that information is valuable over the entire course of the relationship. Expending cost Nc in the first period implies that with probability $p(c)$ the agent learns the entire profile of realizations $(r_{2,t})_{t \geq 1}$. With probability $1 - p(c)$ the agent does not observe any information and there are no more information acquisition opportunities.

Environments \mathcal{P} in this subclass of interest differ by the probability $p(c)$ the agent can learn the profile of returns $(r_{2,t})_{t \geq 1}$. This is equivalent to a per-period expected returns

$r_{\max}(c) = p(c)/3$. Attention is restricted to expected return production functions of the form $r_{\max}(c) = r_{\max}(0) \geq 0$ for $c \in [0, c_{FB}]$ and $r_{\max}(c) = r_{FB}$ for $c \in [c_{FB}, +\infty)$. Furthermore we impose the restriction that $\frac{c_{FB}}{wr_{FB}} = \rho$. In this environment, without loss of efficiency, one can restrict attention to contracts in which aggregate payments Π_N are decided and transferred in the last period. Define $\pi = \Pi_N/N$ the corresponding per-period reward. Reward π need only be conditioned on the following events

- the agent only invests in asset 1 (and obtains returns 0) (event 0)
- the agent invests in asset 2 and only obtains returns 1 when he does (event 1)
- the agent invests in asset 2 and obtains returns -1 at one history (event -1).

It is optimal to discourage the agent to choose asset 2 unless he knows that returns are equal to 1. Indeed, $\pi(-1)$ can be set arbitrarily low and event -1 will not occur on the equilibrium path.

Define $\Delta = \pi(1) - \pi(0)$ the difference in rewards between events 1 and 0. The agent's per-period expected payoff from putting effort c is

$$\begin{aligned} p(c) [(1 - (2/3)^N)\pi(1) + (2/3)^N\pi(0)] + (1 - p(c))\pi(0) - c \\ = \pi(0) + r_{\max}(c) [1 - (2/3)^N] \frac{\Delta}{3} - c, \end{aligned}$$

while the principal's per-period payoff is

$$-\pi(0) + r_{\max}(c) \left(w - [1 - (2/3)^N] \frac{\Delta}{3} \right).$$

Let us first show that any contract such that $\pi(0) \neq 0$ cannot guarantee a positive share of first-best surplus. Indeed, if $\pi(0) < 0$, then for values of r_{FB} low enough, the agent's payoff is strictly negative for all values of $c \in [0, +\infty)$, which implies that the agent doesn't participate and the principal gets profits equal to 0. If instead $\pi(0) > 0$, then for values of r_{FB} low enough, the principal will get negative profits.

Now consider the case where $\pi(0) = 0$. If $\Delta < 3 \frac{c_{FB}}{r_{FB}[1-(2/3)^N]}$, then the agent chooses cost level $c = 0$, which leads to zero profits in environments where $r_{\max}(0) = 0$. Assume now that $\Delta \geq 3 \frac{c_{FB}}{r_{FB}[1-(2/3)^N]}$. For any $\epsilon > 0$, in environments such that $r(0) \frac{\Delta}{3} [1 - (2/3)^N] = r_{FB} \frac{\Delta}{3} [1 - (2/3)^N] - c_{FB} + \epsilon$, the agent chooses to expend cost $c = 0$, and the principal obtains payoff

$$\left(r_{FB} - \frac{c_{FB} - \epsilon}{[1 - (2/3)^N] \Delta/3} \right) (w - [1 - (2/3)^N] \Delta/3)$$

Letting ϵ go to 0 and maximizing over Δ (which gives $\Delta = \frac{3w}{1-(2/3)^N} \sqrt{\rho}$), yields that the principal can guarantee himself a payoff of at most

$$(wr_{FB} - c_{FB}) \left(1 - \frac{2\sqrt{\rho}}{1 + \sqrt{\rho}} \right).$$

This concludes the proof. □

Proof of Fact 2: The fact that benchmark contracts satisfy no-loss is immediate. Let us turn to the converse.

Contract $(\pi_t)_{t \geq 1}$ induces indirect vNM preferences for the agent and the principal over lotteries with outcomes $(w_t, w_t^0)_{t \geq 1}$. Given such a lottery L , the principal and the agent respectively have expected utility

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \quad \text{and} \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right).$$

Because no-loss must hold for every underlying environment \mathcal{P} and every strategy of the agent, it implies that for every probability distribution L over outcomes $(w_t, w_t^0)_{t \geq 1}$,

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t \right) \geq 0 \quad \iff \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) \geq 0.$$

Hence, if $\mathbb{E}_L(\sum_{t=1}^N w_t - w_t^0) = 0$, then $\mathbb{E}_L(\sum_{t=1}^N \pi_t)$ and $-\mathbb{E}_L(\sum_{t=1}^N \pi_t)$ must have the same

sign, which implies that

$$\mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 \right) = 0 \quad \Rightarrow \quad \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) = 0.$$

Consider the deterministic sequence such that for all $t > 1$, $w_t = w_t^0 = 0$, $w_1 = 0$ and $w_1^0 = 1$. Let $\alpha = -\sum_{t=1}^N \pi_t$ for this deterministic sequence of outcomes. Let L_{-1} denote the lottery putting unit mass on this outcome. For any lottery L such that $\mathbb{E}_L(\sum_{t \geq 1} \pi_t) \geq 0$, consider the compound lottery $\hat{L} = pL_{-1} + (1-p)L$, with $p/(1-p) = \mathbb{E}_L(\sum_{t \geq 1} \pi_t)$. By construction, $\mathbb{E}_{\hat{L}}(\sum_{t \geq 1} w_t - w_t^0) = 0$ so that necessarily,

$$\begin{aligned} \mathbb{E}_{\hat{L}} \left(\sum_{t=1}^N \pi_t \right) = 0 &\iff -p\alpha + (1-p)\mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) = 0 \\ &\iff \mathbb{E}_L \left(\sum_{t=1}^N \pi_t \right) = \alpha \mathbb{E}_L \left(\sum_{t=1}^N w_t - w_t^0 \right). \end{aligned}$$

Since this must hold for all lotteries L , it must be that for all t , $\pi_t = \alpha(w_t - w_t^0)$. Finally it is immediate that in order to satisfy no-loss, it must be that $\alpha \in (0, 1)$. \square

B.2 Proofs for Section 4

Proof of Lemma 2: Under any benchmark linear contract, the agent uses conditionally optimal allocation policy a^* . Let (c, a^*) denote the agent's policy under the benchmark contract of parameter α , (\tilde{c}, \tilde{a}) his policy under contract (λ, π) , and (c_0, a^*) the agent's policy in the benchmark contract of parameter α_0 .

By optimality of (\tilde{c}, \tilde{a}) under contract (λ, π) , we have that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left[\Pi_N - \sum_{t=1}^N \tilde{c}_t \right] \geq \mathbb{E}_{c, a^*} \left[\Pi_N - \sum_{t=1}^N c_t \right].$$

We obtain that

$$\mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha S_N - \sum_{t=1}^N \tilde{c}_t \right] + C \geq \mathbb{E}_{c, a^*} \left[\alpha S_N - \sum_{t=1}^N c_t \right] - B \geq \mathbb{E}_{c, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_t \right] - B - \alpha A. \quad (42)$$

By optimality of (c, a^*) under the benchmark contract of parameter α , we have that

$$\mathbb{E}_{c, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_t \right] \geq \mathbb{E}_{c_0, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_{0,t} \right]. \quad (43)$$

By optimality of (c_0, a^*) under the benchmark contract of parameter α_0 we obtain

$$\mathbb{E}_{c_0, a^*} \left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t} \right] \geq \mathbb{E}_{\tilde{c}, a^*} \left[\alpha_0 \Sigma_N - \sum_{t=1}^N \tilde{c}_t \right].$$

Note that by definition of a^* and S_N , $\mathbb{E}_{\tilde{c}, a^*} \Sigma_N \geq \mathbb{E}_{\tilde{c}, \tilde{a}} S_N$. Indeed, under a^* , Σ_T delivers positive expected returns every period, while S_T (under any allocation policy) provides at best a fraction of these returns. This implies that

$$\mathbb{E}_{c_0, a^*} \left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t} \right] \geq \mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha_0 S_N - \sum_{t=1}^N \tilde{c}_t \right]. \quad (44)$$

Combining (42), (43) and (44) yields

$$\begin{aligned} \mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha S_N - \sum_{t=1}^N \tilde{c}_t \right] + \alpha A + B + C &\geq \mathbb{E}_{c_0, a^*} \left[\alpha \Sigma_N - \sum_{t=1}^N c_{0,t} \right] \\ &\geq (\alpha - \alpha_0) \mathbb{E}_{c_0, a^*} \Sigma_N + \mathbb{E}_{c_0, a^*} \left[\alpha_0 \Sigma_N - \sum_{t=1}^N c_{0,t} \right] \\ &\geq (\alpha - \alpha_0) \mathbb{E}_{c_0, a^*} \Sigma_N + \mathbb{E}_{\tilde{c}, \tilde{a}} \left[\alpha_0 S_N - \sum_{t=1}^N \tilde{c}_t \right]. \end{aligned}$$

Altogether, this implies that $(\alpha - \alpha_0) [\mathbb{E}_{c_0, a^*} \Sigma_N - \mathbb{E}_{\tilde{c}, \tilde{a}} S_N] \leq \alpha w \bar{d} (2\sqrt{N} + 1)$. Hence we obtain

that

$$\mathbb{E}_{\bar{c}, \bar{a}}[S_N - \Pi_N] \geq (1 - \alpha)\mathbb{E}_{c_0, \alpha^*}\Sigma_N - (1 - \alpha)\frac{\alpha A + B + C}{\alpha - \alpha_0} - C.$$

Dividing by Nw , this yields that $r_{\lambda, \pi} \geq (1 - \eta)r_{\alpha_0} - \frac{1}{Nw} \left[C + \frac{1 - \eta}{\eta}(\alpha A + B + C) \right]$. \square

B.3 Proofs for Section 5

B.3.1 Proofs for Section 5.1: Calibrating a Broader Class of Contracts

The proof of Theorem 3 uses the following extension of Lemma 1.

Lemma B.1 (incentive approximation). *For any realization of uncertainty,*

$$-\bar{\pi}^\eta \leq \sum_{t=1}^T \pi_t - \pi_t^\eta(\lambda_t) \leq \bar{\pi}^\eta \sqrt{T} \quad (45)$$

$$-\bar{\pi}^\eta \sqrt{T} \leq \max_{T' \leq T} \sum_{t=T'}^T \pi_t^\eta - \pi_t^\eta(\lambda_t) \leq \bar{\pi}^\eta \sqrt{T}. \quad (46)$$

Proof. Let $\mathcal{R}_T = (\mathcal{R}_{1,T}, \mathcal{R}_{2,T})$ denote the vector of regrets and $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$ the vector of flow regrets. Using the fact that $\mathcal{R}_{2,T+1} = \mathcal{R}_{2,T}^+ + (1 - \lambda_T)\pi_{T+1}^\eta$, and exploiting the equality $\mathcal{R}_{2,T}^+(\mathcal{R}_{2,T} - \mathcal{R}_{2,T}^+) = 0$, we have

$$\begin{aligned} \langle \mathcal{R}_T^+, \rho_{T+1} \rangle &= \mathcal{R}_{1,T}^+[\pi_T - \lambda_{T+1}\pi_T^\eta] + \mathcal{R}_{2,T}^+(1 - \lambda_{T+1})\pi_{T+1}^\eta \\ &= \mathcal{R}_{1,T}^+\pi_t + [(1 - \lambda_{T+1})\mathcal{R}_{2,T}^+ - \lambda_{T+1}\mathcal{R}_{1,T}^+]\pi_{T+1}^\eta. \end{aligned}$$

Hence, the contract $(\lambda_t, \pi_t)_{t \geq 1}$ defined by (22) ensures that for all realizations of \mathbf{r}_{T+1} , $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$.

We now prove by induction that $\|\mathcal{R}_T^+\|^2 \leq \sum_{t=1}^T (\pi_t^\eta)^2$. The property clearly holds for $T = 1$. We now assume that it holds at T and show it must hold at $T + 1$. Consider first

the case where $\mathcal{R}_{2,T} > 0$.

$$\begin{aligned} \|\mathcal{R}_{T+1}^+\|^2 &\leq \|\mathcal{R}_T^+ + \rho_{T+1}\|^2 \leq \|\mathcal{R}_T^+\|^2 + 2\langle \mathcal{R}_T^+, \rho_{T+1} \rangle + \|\rho_{T+1}\|^2 \\ &\leq \|\mathcal{R}_T^+\| + \|\rho_{T+1}\|^2 \end{aligned}$$

where we used the fact that by construction, $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$. Furthermore, we have that

$$\begin{aligned} \|\rho_{T+1}\|^2 &\leq (\pi_{T+1} - \pi_{T+1}^\eta(\lambda_{T+1}))^2 + (\mathcal{R}_{2,T}^+ + (1 - \lambda_{T+1})\pi_{T+1}^\eta - \mathcal{R}_{2,T})^2 \\ &\leq \lambda_{T+1}^2(\pi_{T+1}^\eta)^2 + (1 - \lambda_{T+1})^2(\pi_{T+1}^\eta)^2 \\ &\leq (\pi_{T+1}^\eta)^2. \end{aligned}$$

Using the induction hypothesis, this implies that $\|\mathcal{R}_{T+1}\|^2 \leq \sum_{t=1}^{T+1} (\pi_t^\eta)^2$. A similar proof holds when $\mathcal{R}_{2,T} < 0$, taking into account that in this case, $\mathcal{R}_{2,T+1} = (1 - \lambda_{T+1})\pi_{T+1}^\eta$. Hence, by induction, this implies that for all $T \geq 1$, $\|\mathcal{R}_T^+\|^2 \leq \sum_{t=1}^T (\pi_t^\eta)^2$. Since $|\pi_t^\eta| \leq \bar{\pi}^\eta$, this implies inequality (46) and the right-hand side of (45). The left-hand side of (45) follows from an induction identical to that used to prove the left-hand side of (18). \square

Proof of Theorem 3: We begin with the proof of (23). Let (\hat{c}, \hat{a}) denote an optimal strategy for the agent under calibrated contract (λ, π) , and let (c, a) denote an optimal strategy for the agent under benchmark contract $\pi^0 = (\pi_t^0)_{t \geq 1}$. By optimality of \hat{a} under (λ, π) , we obtain that

$$\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \pi_t - \hat{c}_t \right) \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t - c_t \right).$$

By (45) this implies that

$$\mathbb{E}_{\hat{c}, \hat{a}} \left(\sum_{t=1}^N \pi_t^\eta(\lambda_t) - \hat{c}_t \right) + \bar{\pi}^\eta \sqrt{N} \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^\eta(\lambda_t) - c_t \right) - \bar{\pi}^\eta.$$

By (46) we obtain

$$\mathbb{E}_{\widehat{c}, \widehat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^\eta - \widehat{c}_t \right) + \bar{\pi}^\eta \sqrt{N} \geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^\eta - c_t \right) - \bar{\pi}^\eta (1 + \sqrt{N}).$$

Using the fact that (c, a) is optimal under contract $(\pi_t^0)_{t \geq 1}$, and that $\mathbb{E}_{a, c}(\pi_t^0) \geq 0$, this implies that

$$\begin{aligned} \mathbb{E}_{\widehat{c}, \widehat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^0 + \lambda_t \eta (w_t - w_t^0 - \pi_t^\eta) - \widehat{c}_t \right) &\geq \mathbb{E}_{c, a} \left(\sum_{t=1}^N \pi_t^0 + \eta (w_t - w_t^0 - \pi_t^\eta) - c_t \right) - \bar{\pi}^\eta (2\sqrt{N} + 1) \\ &\geq \mathbb{E}_{\widehat{c}, \widehat{a}} \left(\sum_{t=1}^N \lambda_t \pi_t^0 - \widehat{c}_t \right) + \mathbb{E}_{c, a} \left(\sum_{t=1}^N \eta (w_t - w_t^0 - \pi_t^\eta) \right) - \bar{\pi}^\eta (2\sqrt{N} + 1). \end{aligned}$$

Thus, using (45) and the fact that $w_t - w_t^0 - \pi_t^\eta = \frac{1}{1+\eta}(w_t - w_t^0 - \pi_t^0)$, we obtain that

$$\begin{aligned} \mathbb{E}_{\widehat{c}, \widehat{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0 - \pi_t^\eta) \right) - \mathbb{E}_{c, a} \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t^\eta \right) &\geq -\frac{\bar{\pi}^\eta}{\eta} (2\sqrt{N} + 1) \\ \Rightarrow \mathbb{E}_{\widehat{c}, \widehat{a}} \left(\sum_{t=1}^N \lambda_t (w_t - w_t^0) - \pi_t \right) &\geq \frac{1}{1+\eta} \mathbb{E}_{c, a} \left(\sum_{t=1}^N w_t - w_t^0 - \pi_t^0 \right) - \frac{\bar{\pi}^\eta}{\eta} (2\sqrt{N} + 1 + \eta). \end{aligned}$$

Inequality (23) follows from normalizing by $1/wN$. Inequality (24) follows from the fact that by Lemma B.1 incentives are approximately correct from the perspective of any history h_T . Indeed, since payments $(\pi_t)_{t \geq 0}$ are positive, the fact that (45) holds at every time t implies that for all T ,

$$-\bar{\pi}^\eta (\sqrt{N} + 1) \leq \sum_{t=T}^N \pi_t - \pi_t^\eta \leq \bar{\pi}^\eta (\sqrt{N} + 1).$$

Furthermore, (46) implies that

$$-\bar{\pi} \sqrt{N} \leq \sum_{t=T}^N \pi_t^\eta - \pi_t^\eta (\lambda_t) \leq \bar{\pi} \sqrt{N}.$$

Hence a proof identical to that of inequality (23) implies inequality (24). \square

B.3.2 Proofs for Section 5.2: Self Screening by Uninformed Agents

The proof of Lemma 3 requires the following extension of the Azuma-Hoeffding inequality.

Lemma B.2 (an extension of Azuma-Hoeffding). *Consider a martingale with increments Δ_t such that $|\Delta_t| \leq \bar{\gamma}$. Filtration $(\mathcal{F}_t)_{t \geq 1}$ corresponds to the information available at the beginning of period t . Let $\gamma_t \equiv \sup |\Delta_t| | \mathcal{F}_t$ and $T_m \equiv \inf \left\{ T \mid \bar{\gamma}^2 + \sum_{t=1}^T \gamma_t^2 \geq m \right\}$. The following hold.*

$$(i) \quad \forall \kappa > 0, \text{Prob} \left(\sum_{t=1}^{T_m} \Delta_t \geq \kappa \right) \leq \exp \left(-2 \frac{\kappa^2}{m} \right)$$

$$(ii) \quad \forall \kappa > 0, \text{Prob} \left(\max_{T \leq T_m} \sum_{t=1}^T \Delta_t \geq \kappa \right) \leq 2 \exp \left(-2 \frac{\kappa^2}{m} \right).$$

Proof of Lemma B.2: Let us begin with point (i). By Hoeffding's Lemma (see Hoeffding (1963) or Cesa-Bianchi and Lugosi (2006), Lemma 2.2), we have that

$$\mathbb{E}(\exp(\lambda \Delta_t) | \mathcal{F}_t) \leq \exp \left(\frac{\lambda^2 \gamma_t^2}{8} \right).$$

By construction $\sum_{t=1}^{T_m} \gamma_t^2 \leq m$. Hence, using Chernoff's method, we have that for any $\lambda > 0$

$$\begin{aligned} \text{Prob} \left(\sum_{t=1}^{T_m} \Delta_t \geq \kappa \right) &\leq \exp(-\lambda \kappa) \mathbb{E} \left(\prod_{t=1}^{T_m} \exp(\lambda \Delta_t) \right) \\ &\leq \exp(-\lambda \kappa) \mathbb{E} \left(\exp(\lambda \Delta_1) \mathbb{E} \left(\exp(\lambda \Delta_2) \cdots \mathbb{E} \left(\exp(\lambda \Delta_{T_m}) | \mathcal{F}_{T_m} \right) | \cdots | \mathcal{F}_2 \right) \right) \\ &\leq \exp(-\lambda \kappa) \mathbb{E} \left(\exp \left(\frac{\lambda^2}{8} \sum_{t=1}^{T_m} \gamma_t^2 \right) \right) \leq \exp(-\lambda \kappa) \exp \left(\frac{\lambda^2}{8} m \right). \end{aligned}$$

Minimizing over λ (i.e. setting $\lambda = 4\kappa/m$) yields point (i).

Point (ii) follows from point (i) by adapting the standard reflection techniques used for Brownian motions. Let $B_T = \sum_{t=1}^T \Delta_t$. Pick $\kappa > 0$. We want to evaluate $\text{Prob}(\max_{T \leq T_m} B_T \geq \kappa)$. Consider the process $\tilde{B}_T = \sum_{t=1}^T \epsilon_t \Delta_t$, where $\epsilon_t = \mathbf{1}_{[\max_{s < t} B_s] < \kappa} - \mathbf{1}_{[\max_{s < t} B_s] \geq \kappa}$. Process

\tilde{B}_T is a martingale, corresponding to reflecting B_T the first time it crosses level κ . Note also that $|\epsilon_t \Delta_t| = |\Delta_t|$. We have that

$$\begin{aligned} \text{Prob}\left(\max_{T \leq T_m} B_T \geq \kappa\right) &= \text{Prob}(B_{T_m} \geq \kappa) + \text{Prob}(B_{T_m} < \kappa \text{ and } \max_{T \leq T_m} B_T \geq \kappa) \\ &\leq \text{Prob}(B_{T_m} \geq \kappa) + \text{Prob}(\tilde{B}_{T_m} \geq \kappa). \end{aligned} \quad (47)$$

Note that (47) is an inequality, rather than an equality as in the case of a Brownian motion, because of the discreteness of martingale increments. Still this suffices for our purpose. Indeed, by applying point (i) to both B_{T_m} and \tilde{B}_{T_m} , we obtain that indeed, $\text{Prob}\left(\max_{T \leq T_m} \sum_{t=1}^T \Delta_t \geq \kappa\right) \leq 2 \exp\left(-2\frac{\kappa^2}{m}\right)$. This concludes the proof. \square

Proof of Lemma 3: We have that

$$S_T = \sum_{t=1}^T \lambda_t \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0] + \sum_{t=1}^T \lambda_t (w_t - w_t^0 - \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0]).$$

Since the agent is uninformed, by definition of w_t^0 , we have that for all allocation strategies a , $\mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0] \leq 0$. Define $\Delta_t \equiv \lambda_t (w_t - w_t^0 - \mathbb{E}_a[w_t - w_t^0 | \mathcal{F}_t^0]) / w$. Δ_t is a martingale increment such that $|\Delta_t| \leq 2\lambda_t d_t$.

Let us define $\chi_T = \bar{d}^2 + \sum_{t=1}^T \lambda_t^2 d_t^2$. For all $m \in \mathbb{N}$, let T_m denote the stopping time $\inf\{T | \chi_T \geq m\}$. Using Lemma B.2, we obtain that for all m

$$\begin{aligned} \text{Prob}\left(S_{T_m} \geq 2w\sqrt{\chi_{T_m}}\sqrt{M + \ln \chi_{T_m}}\right) &\leq \text{Prob}\left(\sum_{t=1}^{T_m} \Delta_t \geq 2\sqrt{\chi_{T_m}}\sqrt{M + \ln \chi_{T_m}}\right) \\ &\leq \exp(-2(\ln m + M)) \leq \exp(-2M) \frac{1}{m^2}. \end{aligned}$$

In addition, conditional on $S_{T_m} \leq 2w\sqrt{\chi_{T_m}}\sqrt{M + \ln \chi_{T_m}}$, Lemma B.2 implies that the

probability that there exists $T \in [T_m, T_{m+1} - 1]$ such that $S_T \geq \Theta_T$ is less than

$$\text{Prob} \left(\sup_{T \in \{T_m, \dots, T_{m+1}-1\}} \sum_{t=T_m}^T \Delta_t \geq 2\sqrt{M + \ln m} \right) \leq 2 \exp(-2M) \frac{1}{m^2}.$$

Hence it follows that

$$\mathbb{E}_a \left(\sum_{t=1}^N \mathbf{1}_{S_t > \Theta_t} \right) \leq 3 \exp(-2M) \sum_{m \in \mathbb{N}} \frac{1}{m^2} \leq \frac{\pi^2}{2} \exp(-2M).$$

This concludes the proof. \square

Let us now turn to the proof of Theorem 4. Let $\Pi_T^\Theta = \sum_{t=1}^T \pi_t^\Theta$ denote actual rewards, up to time T . The following lemma extends Lemma 1.

Lemma B.3 (approximate incentives). *For all $T, T' < T$, and all paths of play, we have that*

$$\Sigma_{T \setminus T'} - S_{T \setminus T'} \leq w \sqrt{\sum_{t=1}^T d_t^2} \quad (48)$$

$$-\alpha \Theta_T - \alpha w \bar{d} - b \leq \Pi_T^\Theta - \alpha S_T \leq \alpha w \sqrt{\sum_{t=1}^T d_t^2}. \quad (49)$$

Proof. A proof identical to that of Lemma 1 yields the left-hand side of (48) and the right-hand side of (49). The left-hand side of (49) is proven by induction. Assume it holds at time T . If $\alpha S_{T+1} - \alpha \Theta_{T+1} \leq 0$, then the inequality holds trivially. Consider now the case where $\alpha S_{T+1} - \alpha \Theta_{T+1} > 0$. If $\Pi_T \geq \alpha S_T - \alpha \Theta_T$ then we necessarily have $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha \Theta_{T+1} - \alpha w \bar{d}$ since Θ_T is increasing in T . If instead, $\Pi_T \in [\alpha S_T - \alpha \Theta_T - \alpha w \bar{d}, \alpha S_T - \alpha \Theta_T]$, then necessarily, $\Pi_T < \alpha S_T$, so that $\lambda_{T+1} = 1$ and $\pi_{T+1} = \alpha(w_{T+1} - w_{T+1}^0)^+$. It follows that $\Pi_{T+1} \geq \alpha S_{T+1} - \alpha \Theta_{T+1} - \alpha w \bar{d}$. \square

Proof of Theorem 4: Combining Lemmas 2, 3 and B.3 yields Theorem 4. \square

It is worth noting that the efficiency bound given in Theorem 4 can be improved in some circumstances. The following theorem shows that when returns are either zero or bounded away from 0, the performance loss will be of order $\sqrt{1/N}$ rather than $\sqrt{\ln N/N}$.

Assumption 3 (grainy returns). *Let (c, a^*) denote the agent's policy under the benchmark contract with rate α_0 . There exists $\xi > 0$ such that whenever $\mathbb{E}_{c, a^*}[w_t - w_t^0 | \mathcal{F}_t] > 0$, then $\mathbb{E}_{c, a^*}[w_t - w_t^0 | \mathcal{F}_t] > \xi$.*

Theorem B.1. *Pick α_0 and for any $\eta > 0$, set $\alpha = \alpha_0 + \eta(1 - \alpha_0)$. If Assumption 3 holds, there exists a constant m such that for all N and all probability spaces \mathcal{P} ,*

$$r_{\lambda, \pi^\Theta} \geq (1 - \eta)r_{\alpha_0} - m \frac{1}{\sqrt{N}}.$$

Proof. As in the case of Theorem 4, the proof strategy is to adapt the the bounds of Lemma 1 and apply the reasoning of Lemma 2. Specifically, the left-hand side of bound (49) in Lemma B.1 must be improved. Let (c, a^*) denote the agent's optimal strategy under the benchmark contract of parameter α . To exploit the reasoning of Lemma 2 and obtain Theorem B.1, it is sufficient to prove a bound of the form

$$-B \leq \mathbb{E}_{c, a^*} [\Pi_N^\Theta - \alpha S_N], \quad (50)$$

where B is a number independent of N and \mathcal{P} . We show that this is indeed the case. By construction, we have that

$$\mathbb{E}_{c, a^*} \Pi_N^\Theta \geq \mathbb{E}_{c, a^*} \alpha S_N - \alpha w \bar{d} - \alpha w \bar{d} \mathbb{E}_{c, a^*} \left[\sum_{T=1}^N \mathbf{1}_{S_T < \Theta_T} \right].$$

Hence, it is sufficient to show that under (c, a^*) , the expected number of periods where the hurdle is not met is bounded above by a constant independent of N .

Let $\Delta_t = w_t - w_t^0 - \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t]$ and $\chi_T = \bar{d}^2 + \sum_{t=1}^T d_t^2$. Note that under strategy (c, a^*) ,

Assumption 3 implies that if $d_t > 0$, then $\mathbb{E}_{c,a^*}(w_t - w_t^0 | \mathcal{F}_t) > \xi$. Hence $\sum_{t=1}^T \mathbb{E}_{c,a^*}(w_t - w_t^0 | \mathcal{F}_t) \geq \xi(\chi_T/\bar{d}^2 - 1)$. By (48), for any T ,

$$\begin{aligned} \text{Prob}_{c,a^*}(S_T < \Theta_T) &\leq \text{Prob}_{c,a^*}(\Sigma_T < \Theta_T + w\sqrt{\chi_T}) \\ &\leq \text{Prob}_{c,a^*}\left(\sum_{t=1}^T \mathbb{E}[w_t - w_t^0 | \mathcal{F}_t] + \sum_{t=1}^T \Delta_t < \Theta_T + w\sqrt{\chi_T}\right) \\ &\leq \text{Prob}_{c,a^*}\left(\xi \left[\frac{\chi_T}{\bar{d}^2} - 1\right] + \sum_{t=1}^T \Delta_t < \Theta_T + w\sqrt{\chi_T}\right) \\ &\leq \text{Prob}_{c,a^*}\left(\sum_{t=1}^T \Delta_t < -\xi \left[\frac{\chi_T}{\bar{d}^2} - 1\right] + \Theta_T + w\sqrt{\chi_T}\right). \end{aligned}$$

An argument similar to that used in the proof of Lemma 3 yields that

$\sum_{T=1}^{+\infty} \text{Prob}\left(\sum_{t=1}^T \Delta_t < -\frac{\xi}{\bar{d}^2}\chi_T + \xi + \Theta_T + w\sqrt{\chi_T}\right)$ is bounded above by a constant. This concludes the proof. \square

B.3.3 Proofs for Section 5.3 (multi-agent contracts)

Proof of Lemma 4: Optimal strategies for the agents $(c^*, a^*) = (c_k^*, a_k^*)_{k \in K}$ are such that for any other profile of strategies $(c, a) = (c_k, a_k)_{k \in K}$, and for all $k \in K$,

$$\mathbb{E}_{c_k^*, a_k^*} \left[\sum_{t=1}^N \alpha(w_{k,t} - w_{k-1,t}) - c_{k,t}^* \right] \geq \mathbb{E}_{c_k, a_k} \left[\sum_{t=1}^N \alpha(w_{k,t} - w_{k-1,t}) - c_{k,t} \right].$$

Summing over k , this implies that

$$\mathbb{E}_{c^*, a^*} \left[\sum_{t=1}^N \alpha(w_{K,t} - w_{0,t}) - \sum_{k \in K} c_{k,t}^* \right] \geq \mathbb{E}_{c, a} \left[\sum_{t=1}^N \alpha(w_{K,t} - w_{0,t}) - \sum_{k \in K} c_{k,t} \right].$$

This implies that

$$\mathbb{E}_{c^*, a^*} \left[\sum_{t=1}^N (1 - \alpha)(w_{K,t} - w_{0,t}) \right] \geq \frac{1 - \alpha}{\alpha} \mathbb{E}_{c, a} \left[\sum_{t=1}^N \alpha(w_{K,t} - w_{0,t}) - \sum_{k \in K} c_{k,t} \right].$$

Since this holds for any strategy profile (c, a) , we obtain that indeed

$$r_\alpha \geq (1 - \alpha) \max_{c \in [0, +\infty)} \left(r_{\max}(c) - \frac{c}{\alpha w} \right).$$

□

Proof of Theorem 5: The proof is similar to that of Theorem 2. Let $\mathcal{R}_T = (\mathcal{R}_{k,T}^1, \alpha \mathcal{R}_{k,T}^2)_{k \in \{1, \dots, K\}}$ denote the vector of regrets. We first show that regrets are at most of order \sqrt{KN} . Let $\rho_T = \mathcal{R}_T - \mathcal{R}_{T-1}$ and consider the dot product

$$\begin{aligned} \langle \mathcal{R}_T^+, \rho_{T+1} \rangle &= \sum_{k \in K} [\mathcal{R}_{k,T}^1]^+ [\pi_{k,T+1} - \alpha(w_{k,T+1}^\lambda - w_{k-1,T+1}^\lambda)] + \alpha^2 [\mathcal{R}_{k,T}^2]^+ (w_{k,T+1} - w_{k,T+1}^\lambda) \\ &= \sum_{k \in K} [\mathcal{R}_{k,T}^1]^+ \pi_{k,T+1} + \alpha \left(\alpha [\mathcal{R}_{k,T}^2]^+ (1 - \lambda_{k,T+1}) - [\mathcal{R}_{k,T}^1]^+ \lambda_{k,T+1} \right) (w_{k,T+1} - w_{k-1,T+1}^\lambda). \end{aligned}$$

This implies that contract parameters (λ_t, π_t) are such that $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$. Using a proof identical to that of Lemma 1, the fact that $\langle \mathcal{R}_T^+, \rho_{T+1} \rangle = 0$ and the fact that $\|\rho_{T+1}\|^2$ is bounded above by $Kw^2\bar{d}^2$ imply that for every k ,

$$\begin{aligned} \sum_{t=1}^N \pi_{k,t} - \alpha(w_{k,t}^\lambda - w_{k-1,t}^\lambda) &\leq w\bar{d}\sqrt{KN} \\ \max_{T' \leq T} \sum_{t=T'}^N w_{k,t} - w_{k,t}^\lambda &\leq w\bar{d}\sqrt{KN}. \end{aligned}$$

Furthermore, by construction, we have that

$$\sum_{t=1}^N \pi_{k,t} - \alpha(w_{k,t}^\lambda - w_{k-1,t}^\lambda) \geq -w\bar{d}.$$

Hence the conditions to apply Lemma 2 hold, and it follows that

$$r_{\pi, \lambda} \geq (1 - \eta)r_{\alpha_0} - m\sqrt{\frac{K}{N}}$$

For m independent of N , K and \mathcal{P} . As in the case of Theorem 2, a similar proof holds from the perspective of every history h_T . □

References

- AL-NAJJAR, N. AND J. WEINSTEIN (2008): “Comparative Testing of Experts,” *Econometrica*, 76, 541–559.
- BATTAGLINI, M. (2005): “Long-term contracting with markovian consumers,” *The American economic review*, 95, 637–658.
- BERGEMANN, D. AND S. MORRIS (2005): “Robust mechanism design,” *Econometrica*, 73, 1771–1813.
- BERNHEIM, B. D. AND D. RAY (1989): “Collective dynamic consistency in repeated games,” *Games and Economic Behavior*, 1, 295–326.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J. ROCHET (2007): “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications,” *Review of Economic Studies*, 74, 345–390.
- BIAIS, B., T. MARIOTTI, J. ROCHET, AND S. VILLENEUVE (2010): “Large risks, limited liability, and dynamic moral hazard,” *Econometrica*, 78, 73–118.
- BILLINGSLEY, P. (1995): *Probability and Measure*, John Wiley & Sons, New York.
- BLACKWELL, D. (1956): “An Analog of the Minimax Theorem for Vector Payoffs,” *Pacific Journal of Mathematics*, 6, 1–8.
- BOSE, A., D. PAL, AND D. SAPPINGTON (2011): “On the performance of linear contracts,” *Journal of Economics & Management Strategy*, 20, 159–193.
- CESA-BIANCHI, N. AND G. LUGOSI (2006): *Prediction, Learning, and Games*, Cambridge University Press.
- CHASSANG, S. (2011): “Calibrated Incentive Contracts,” *Princeton University Economic Theory Center Working Paper No. 013-2011*.

- CHU, L. AND D. SAPPINGTON (2007): “Simple cost-sharing contracts,” *The American Economic Review*, 97, 419–428.
- CHUNG, K. AND J. ELY (2003): “Implementation with Near-Complete Information,” *Econometrica*, 71, 857–871.
- CHUNG, K. S. AND J. ELY (2007): “Foundations of Dominant-Strategy Mechanisms,” *Review of Economic Studies*, 74, 447–476.
- DASGUPTA, P., P. HAMMOND, AND E. MASKIN (1979): “The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility,” *The Review of Economic Studies*, 46, 185–216.
- DEMARZO, P. AND M. FISHMAN (2007): “Optimal Long-Term Financial Contracting,” *Review of Financial Studies*, 20, 2079.
- DEMARZO, P., I. KREMER, AND Y. MANSOUR (2009): “Hannan and Blackwell meet Black and Scholes: Approachability and Robust Option Pricing,” *Working paper, Stanford University*.
- DEMARZO, P. AND Y. SANNIKOV (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *The Journal of Finance*, 61, 2681–2724.
- ECHENIQUE, F. AND E. SHMAYA (2007): “You Won’t Harm Me if You Fool Me,” *Working Paper, Caltech*.
- EDMANS, A., X. GABAIX, T. SADZIK, AND Y. SANNIKOV (forthcoming): “Dynamic Incentive Accounts,” *Journal of Finance*.
- ELIAZ, K. (2002): “Fault Tolerant Implementation,” *Review of Economic Studies*, 69, 589–610.
- ESCOBAR, J. AND J. TOIKKA (2009): “A Folk Theorem with Markovian Private Information,” *Working Paper, Massachusetts Institute of Technology*.
- FARRELL, J. AND E. MASKIN (1989): “Renegotiation in repeated games,” *Games and economic behavior*, 1, 327–360.
- FEINBERG, Y. AND C. STEWART (2008): “Testing Multiple Forecasters,” *Econometrica*, 76, 561–582.

- FONG, K. (2008): “Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons,” *Working Paper, Stanford University*.
- FOSTER, D. AND R. VOHRA (1998): “Asymptotic Calibration,” *Biometrika*, 85, 379.
- (1999): “Regret in the Online Decision Problem,” *Games and Economic Behavior*, 29, 7–35.
- FOSTER, D. AND P. YOUNG (2010): “Gaming Performance Fees By Portfolio Managers,” *The Quarterly Journal of Economics*.
- FUDENBERG, D. AND D. LEVINE (1999): “An Easier Way to Calibrate,” *Games and economic behavior*, 29, 131–137.
- GARRETT, D. AND A. PAVAN (2010): “Dynamic Managerial Compensation: On the Optimality of Seniority-based Schemes,” .
- GOETZMANN, W., J. INGERSOLL, AND S. ROSS (2003): “High-Water Marks and Hedge Fund Management Contracts,” *The Journal of Finance*, 58, 1685–1718.
- GOETZMANN, W., J. INGERSOLL, M. SPIEGEL, AND I. WELCH (2007): “Portfolio Performance Manipulation and Manipulation-Proof Performance Measures,” *Review of Financial Studies*.
- GRADWOHL, R. AND Y. SALANT (2011): “How to Buy Advice,” *Working Paper, Northwestern University*.
- GREEN, E. (1987): “Lending and the smoothing of uninsurable income,” in *Contractual arrangements for intertemporal trade*, ed. by E. Prescott and N. Wallace, University of Minnesota Press, 3–25.
- HAGERTY, K. AND W. ROGERSON (1987): “Robust Trading Mechanisms,” *Journal of Economic Theory*, 42, 94–107.
- HANNAN, J. (1957): “Approximation to Bayes Risk in Repeated Play,” in *Contributions to the Theory of Games*, ed. by M. Dresher, A. Tucker, and P. Wolfe, vol. 3, 97–139.
- HARTLINE, J. D. AND T. ROUGHGARDEN (2008): “Optimal Mechanism Design and Money Burning,” in *Symposium on Theory Of Computing (STOC)*, 75–84.

- HE, Z. (2009): “Optimal executive compensation when firm size follows geometric brownian motion,” *Review of Financial Studies*, 22, 859.
- HE, Z. AND W. XIONG (2010): “Delegated Asset Management and Investment Mandates,” *Working Paper, Princeton University*.
- HOEFFDING, W. (1963): “Probability Inequalities for Sums of Bounded Random Variables,” *Journal of the American Statistical Association*, 58, 13–30.
- HOLMSTRÖM, B. AND P. MILGROM (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55, 303–328.
- HURWICZ, L. AND L. SHAPIRO (1978): “Incentive structures maximizing residual gain under incomplete information,” *The Bell Journal of Economics*, 9, 180–191.
- JACKSON, M. AND H. SONNENSCHNEIN (2007): “Overcoming Incentive Constraints by Linking Decisions,” *Econometrica*, 75, 241–257.
- LAFFONT, J. AND J. TIROLE (1988): “The dynamics of incentive contracts,” *Econometrica*, 1153–1175.
- LEHRER, E. (2001): “Any Inspection is Manipulable,” *Econometrica*, 69, 1333–1347.
- LO, A. (2001): “Risk Management for Hedge Funds: Introduction and Overview,” *Financial Analysts Journal*, 57, 16–33.
- MADARÁSZ, K. AND A. PRAT (2010): “Screening with an Approximate Type Space,” *Working Paper, London School of Economics*.
- NEEMAN, Z. (2003): “The effectiveness of English auctions,” *Games and Economic Behavior*, 43, 214–238.
- OLSZEWSKI, W. AND M. PESKI (2011): “The principal-agent approach to testing experts,” *American Economic Journal: Microeconomics*, 3, 89–113.
- OLSZEWSKI, W. AND A. SANDRONI (2008): “Manipulability of Future-Independent Tests,” *Econometrica*, 76, 1437–1466.
- PANAGEAS, S. AND M. WESTERFIELD (2009): “High-Water Marks: High Risk Appetites? Convex Compensation, Long Horizons, and Portfolio Choice,” *The Journal of Finance*, 64, 1–36.

- PAVAN, A., I. SEGAL, AND J. TOIKKA (2010): “Infinite-Horizon Mechanism Design,” *Working Paper, Northwestern University*.
- RADNER, R. (1981): “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, 49, 1127–1148.
- (1985): “Repeated Principal-Agent Games with Discounting,” *Econometrica*, 53, 1173–1198.
- ROGERSON, W. (1985): “Repeated Moral Hazard,” *Econometrica*, 53, 69–76.
- (2003): “Simple menus of contracts in cost-based procurement and regulation,” *The American economic review*, 93, 919–926.
- RUBINSTEIN, A. (1979): “An optimal conviction policy for offenses that may have been committed by accident,” in *Applied game theory*, ed. by J. Brams, A. Schotter, and G. Schwodiauer, Physica-Verlag, Wurzburg, 406–413.
- RUBINSTEIN, A. AND M. YAARI (1983): “Repeated insurance contracts and moral hazard,” *Journal of Economic Theory*, 30, 74–97.
- SANNIKOV, Y. (2007): “Agency problems, screening and increasing credit lines,” *Working paper, Princeton University*.
- (2008): “A Continuous-Time Version of the Principal–Agent Problem,” *Review of Economic Studies*, 75, 957–984.
- SPEAR, S. AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *The Review of Economic Studies*, 54, 599–617.
- TCHISTYI, A. (2006): “Security design with correlated hidden cash flows: The optimality of performance pricing,” *Working Paper, University of California at Berkeley*.
- WILSON, R. B. (1987): “Game-Theoretic Analyses of Trading Processes,” in *Advances in Economic Theory: Fifth World Congress*, ed. by T. Bewley, Cambridge UK: Cambridge University Press, chap. 2, 33–70.
- ZHU, J. (2010): “Sticky Incentives and Dynamic Agency,” *Working Paper, University of California at Berkeley*.